Deciding on Personalized Ads: Nudging Developers About User Privacy

Mohammad Tahaei
School of Informatics
University of Edinburgh

Alisa Frik
ICSI
University of California, Berkeley

Kami Vaniea
School of Informatics
University of Edinburgh

Abstract

Mobile advertising networks present personalized advertisements to developers as a way to increase revenue. These types of ads use data about users to select potentially more relevant content. However, choice framing also impacts app developers' decisions which in turn impacts their users' privacy. Currently, ad networks provide choices in developer-facing dashboards that control the types of information collected by the ad network as well as how users will be asked for consent. Framing and nudging have been shown to impact users' choices about privacy, we anticipate that they have a similar impact on choices made by developers. We conducted a survey-based online experiment with 400 participants with experience in mobile app development. Across six conditions, we varied the choice framing of options around ad personalization. Participants in the condition where privacy consequences of ads personalization are highlighted in the options are significantly (11.06 times) more likely to choose non-personalized ads compared to participants in the Control condition with no information about privacy. Participants' choice of ad type is driven by impact on revenue, user privacy, and relevance to users. Our findings suggest that developers are impacted by interfaces and need transparent options.

1 Introduction

Mobile advertising networks play an intermediary role of matching advertisers (companies that want to advertise their products) with publishers (apps that want to generate revenue by hosting advertising). They are a popular monetisa-

tion approach [11, 47, 57, 93, 107], with about 77% of free Android apps containing an ad library [48, 51]. To show personalized ads, ad networks collect data from app users, which raises privacy concerns [41, 111, 116]. Targeted ads can also seem intrusive and discriminating to some users [63, 83, 88, 117]. Major operating systems give users an option to limit these ads and associated tracking. However, behavioral research shows that due to status quo bias, people rarely change the default configurations [3, 52, 87, 91], and poor usability makes it hard for users to opt out of behavioral advertising and tracking [45, 56, 90]. Thus, developers' decisions regarding the defaults for their apps have implications for user privacy. Specifically, when configuring ad networks, developers can choose in the developer dashboard between personalized and non-personalized ads. Here again, status quo bias may not play out in favor of user privacy: if ad networks set personalized ads that imply more extensive personal data collection as default choices, it might nudge developers to stick to those privacy-unfriendly defaults [33, 68].

With about 24 million software developers (estimated to go up to 28.7 million by 2024) [82], who are in charge of building apps for personal smart devices, cars, and large industries, it is essential to understand how services they use may impact their decisions. Indeed, studies of privacy-related questions on Stack Overflow [106] and Reddit Android forums [59] show that developers' privacy concerns are heavily driven by large platforms such as Google and Apple. Moreover, there is a growing use of dark patterns that persuade users into make decisions that are in favor of platforms; for example, by using preselected default options, or sneaking a small product or service into users shopping basket without informing users, such as adding travel insurance during the plane ticket purchasing [42, 65, 79]. The use of dark patterns in the context of software development may have negative implications for users, as developers' choices will effect all users of their apps. For example, collecting location data, showing unrestricted ads categories, and displaying personalized ads are often allowed by default in popular ad networks [68, 102].

Similarly, given that ads tailored to users' preferences have a higher value [64], ad networks have incentive to nudge developers into choosing personalized ads over non-personalized ones, without necessarily acknowledging the trade-offs between revenue, user privacy, and experience. In addition to status quo bias leveraged by default choices, salience effect can be leveraged to further facilitate the nudging [18, 92]. For example, while an emphasis on user privacy may steer developers' decisions towards non-personalized ads, an emphasis on potentially larger revenue may nudge developers to choose personalized ads which is used by some ad networks through including statements like "including personalized ads may likely result in higher revenue" in their documentation, quick start guides, and blog posts [102, 104, 108].

In this study, we aim to understand how choice framing in ad networks effects developers' decision making. Our research question are:

RQ1: How does choice framing in ad networks impact developers' decisions about ad personalization?

RQ2: What are the reasons behind developers' choices of personalized or non-personalized ads?

To answer our research questions, we conducted an online survey-based experiment with 400 participants with app development experience. In a hypothetical scenario, we asked them to make a series of choices to integrate ads in a personal finance management app and a gaming app. The main decision of interest was regarding the choice between personalized and non-personalized ads. The framing of those choices was manipulated between one control and five experimental conditions, to emphasize implications for framing around data processing restrictions, user-facing descriptions, user privacy, developer's revenue, and both user privacy and developer's revenue. To help further contextualize and interpret the results, we also surveyed participants' opinions and attitudes about personalized ads, ad networks, and privacy regulations.

We find that although on average the majority of participants decided to integrate the personalized ads, choice framing significantly impacted their decisions. When user privacy implications were made salient, participants were 11.06 times more likely to select non-personalized ads than when the neutral framing was used (Control condition). When a framing emphasized data processing restrictions, participants were 3.45 more likely to select the non-personalized ads than in the Control condition. Other nudges—emphasizing the consequences of ads on an app's revenue, presenting participants with an explicit choice between user privacy and app's revenue, and telling participants that users will be able to see whether the app is using ads based on their personal data or not—did not significantly change participants decisions compared to the Control condition.

The analysis of open-ended responses revealed a variety of reasons for developers' choices, ranging from maximizing the app's revenue and relevance of ads to the uses, to concerns about user privacy and regulation compliance, and implica-

tions for user experience. From the exit survey, we found that even when upper and middle management choose the ad networks and app's business models, developers still feel involved in this decision-making process. However, developers generally believe that they do not have full control over ad networks' data collection, and believe users have even less control. By illustrating the potential impact of choice framing on ad personalization decisions during app development, our results inform regulators about the need to enforce greater control over ad networks' data collection and analysis practices, discourage from using dark patterns, and encourage ad networks to adopt interfaces for developers that may assist them in making informed decisions about user privacy.

2 Related Work

Ad Networks. Ad networks are a popular mobile app monetisation approach [11, 47, 57, 93, 107]. Over half of Android apps include ad network libraries [11, 48, 51, 107], which often offer both personalized and non-personalized ads. Personalized ads attract more user attention than non-personalized ads [20, 63], generating higher engagement and therefore revenue. To provide ads tailored to a specific user, ad networks collect personal information from users such as age, gender, and location [84, 100], not only in free apps that rely mostly on ads to generate revenue, but also in paid apps [19, 47]. However, personalized ads have some negative consequences for users. For example, some users find them discomforting [63, 117], discriminating [86], and intrusive [83, 88].

Options Provided by Ad Networks to Users and Developers. Both users and developers can limit data collection and turn off ad personalization. After the introduction of the General Data Protection Regulation (GDPR) [39] and the California Consumer Privacy Act (CCPA) [26], the prevalence of these options particularly increased [50].

On the user side, self-regulatory programs (e.g., Digital Advertising Alliance opt-out [31]), smartphone operating systems, service providers, and browsers offer settings that allow opting out of ad personalization [66], and at minimum, request user consent to show personalized ads. Research shows limited effectiveness, usefulness, legal compliance [37, 46, 67, 112], and usability [67, 81] of these methods.

On the developer side, ad networks provide an interface for configuring personalization and data collection for specific apps and geographic regions. These interfaces often use defaults that are not in favor of user privacy [68, 102]. Developers tend to keep the defaults, follow industry standards, guidelines, and requirements provided by the platforms built by large tech companies [43, 59, 95, 106] without fully considering all the options and consequences of their choices on user privacy [30, 33, 68]. Developers generally acknowledge the value of user privacy [33, 68, 94], but find it challenging

to understand what information is collected, how it is used by platforms [33, 68, 103], and how to protect user privacy [59, 106]. Hence, some poor user privacy elements in how apps integrate ad networks may be caused by the way ad networks are framing choices and nudging developers through defaults.

Nudging. Humans can be nudged towards making certain actions through the use of specific wordings, framing, colors, and default values [3, 27]. *Choice framing*, in particular, uses the activation of salience effects [18, 92] and status quo bias [52, 87, 91], to effectively nudge the privacy choices of users [3, 16]. For example, priming survey respondents about privacy using words like "privacy-sensitive" and "potential privacy risks" increases the reported privacy concerns [25] and making privacy information salient drives more privacy-preserving choices in user experiments [109]. We believe that similar effects can be achieved in the context of software development, where choice framing in tools and interfaces may affect developers' decision making.

Nudges can be used to encourage users to make decisions that are favorable to service providers (e.g., ad networks) but not necessarily favorable to themselves. Such practices are often referred to *dark patterns*—"instances where designers use their knowledge of human behavior (e.g., psychology) and the desires of users to implement deceptive functionality that is not in the user's best interest" [42, p. 1]. In the context of privacy, the examples of dark patterns include privacy consent forms that do not provide a "reject all" button [81] and hard-to-find (or completely absent) options for deleting accounts [23]. Similar patterns are also visible in ad networks' developer dashboards where the default values are all set to personalized ads and location data is often collected by default [68, 102].

Our Contribution. We extend the literature on developer-facing privacy interfaces by looking at the privacy nudges directed at developers and exploring the impact of choice framing in ad networks' developer dashboards.

3 Method

To answer our research questions, we conducted an online survey-based between-subject experiment with 400 participants with mobile development experience administered using Qualtrics. The study received ethical approval from our institute. All participants provided informed consent before completing the study. We describe the study protocol below, and the full survey text is in Appendix A.2.

After screening for app development experience (Section 3.2), participants were randomly assigned to one of six conditions (Section 3.1), and asked to complete the main survey. Each participant was presented with two hypothetical scenarios in a random order: one was about a *gaming app*, another one was about a *financial app* for personal finance

management. We chose these app categories, because personal finance management has obvious privacy implications (e.g., developers reported more sensitive variables for the financial category compared to other app categories [17]), and gaming is the most popular category on both Apple App Store and Google Play [76, 77].

Participants were asked to imagine that they were a shareholder in a software development company, and together with a small team, they created a (financial or gaming) app, which will be published in Europe and the United States and is mainly targeted towards adults above the age of 18. Then, we asked them to answer questions posed by the "Acme Assistant", a tool for an imaginary ad network that helps with integrating the ad network into the app. The Assistant was inspired by MoPub Integration Suite, a new service by Twitter's MoPub ad network for an easy app integration [74]. The Assistant asked five multiple-choice questions about ad formats (e.g., banner and interstitial), level of graphics (high-quality and moderate-quality), platforms (e.g., Android and iOS), types of ads (personalized and non-personalized), and the regulations that apply to the app (e.g., GDPR, CCPA). After making the choices, they were also asked an open-ended question about the primary reason for choosing the personalized or non-personalized ad type.

After completing the above for both the financial and gaming apps, they were sent to an exit survey with the questions about: how they would go about asking for user consent for the personalized ads, how the choice of ad type would affect an app's revenue or number of users, what role does user privacy play in their daily development routines, and how much users and developers have control over data collected by ad networks. The exit survey provided additional insights about participants' opinions, knowledge, and attitudes, and helped to further contextualize and interpret experimental results. Finally, they answered software and mobile development, and demographics questions.

3.1 Experimental Conditions

All participants were randomly assigned to one of six conditions including one Control group and five treatment groups. The only difference among the conditions was the framing of the choice about personalized or non-personalized ads. The order of all options was randomized. Each choice consisted of a short label phrase followed by a longer description.

Control–Minimal Information ($N = 66$): *(1) Personalized ads: Acme can show personalized ads to your users. (2) Non-personalized ads: Acme will show only non-personalized ads to your users.* This framing was inspired by Google AdMob's developer dashboard to help developers build GDPR-compliant apps for European users (Figure 2 in the Appendix). It used neutral wording about ad types without mentioning any information about collection and processing of user data.

Data Processing Restrictions ($N = 67$): *(1) Ads with unrestricted data processing: Acme can show personalized ads to your users based on a user's past behavior, such as previous visits to sites or apps or where the user has been. (2) Ads with restricted data processing: Acme will show only non-personalized ads to your users based on contextual information, such as the content of your site or app, restricting the use of certain unique identifiers and other data.* This framing was inspired by Google AdMob's developer dashboard to help developers build CCPA-compliant apps for California users (Figure 3 in the Appendix) and it explicitly hinted at the types of data used for ad personalization, which may indirectly encouraged developers to consider privacy implications of such data processing. We based two of our conditions on Google AdMob because it is the most common mobile ad network in apps [6, 7, 40].

User-Facing Descriptions ($N = 68$): *(1) Ads with 'Personalized Ads' tag displayed to users: Acme can show personalized ads to your users. Users will see the 'Personalized Ads' tag next to the 'Install' button and the following text in your app description in the App Store or Google play "This app shows ads personalized based on your personal information." (2) Ads with 'Non-personalized Ads' tag displayed to users: Acme will show only non-personalized ads to your users. Users will see the 'Non-personalized Ads' tag next to the 'Install' button and the following text in your app description in the App Store or Google play "This app shows ads not personalized based on your personal information."* This condition aimed at leveraging transparency and nudging developers' accountability and responsibility to users. The framing was inspired by the recent additions to the Apple App Store called "Privacy Details" to "help users better understand an app's privacy practices before they download the app on any Apple platform" [12] and prior work's recommendation about including privacy features of apps in the app stores to softly nudge developers to consider user privacy in their apps [59].

Privacy Focused ($N = 67$): *(1) Ads with lower user privacy: Acme can show personalized ads to your users based on their past behavior, such as previous visits to sites or apps or where the user has been. (2) Ads with higher user privacy: Acme will show only non-personalized ads to your users based on contextual information, such as the content of your site or app.* This condition is aimed at leveraging salience effects [18, 92], by making privacy implications prominent in the choice option descriptions.

Revenue Focused ($N = 65$): *(1) Ads with higher revenue: Acme can show personalized ads to your users, which may yield higher revenue than non-personalized ads. (2) Ads with lower revenue: Acme will show only non-personalized ads to your users, which may yield lower revenue than personalized ads.* This condition aimed at leveraging salience effects [18, 92], by making revenue implications prominent in the choice option descriptions.

Privacy vs. Revenue ($N = 67$): *(1) Ads with higher revenue: Acme can show personalized ads to your users, which may yield higher revenue than non-personalized ads. (2) Ads with higher user privacy: Acme will show only non-personalized ads to your users which may increase your users' privacy.* This condition aimed at exploring what choices the participants would make if they were faced with an explicit trade-off between the user privacy and revenue.

3.2 Recruitment and Screening

We used Prolific, GitHub, and LinkedIn groups to recruit the participants (Jan '21). On average, the survey took 19 minutes ($SD = 89$, *median* $= 13$) to complete. The large standard deviation was due to some participants who left the survey open but stepped away before returning and completing it.

Prolific. Using Prolific's exclusion criteria, we recruited 1,288 participants who were fluent in English, had computer programming skills, and an approval rate of at least 90%. They responded to a 1-minute screening survey (Appendix A.1) to assess their software development experience, and received £0.15 compensation. Those who worked on at least one app in the past three years ($N = 466$) were invited to the main survey and were paid £1.50 for completing it. Of the invited participants, 372 respondents started the main survey, but eight did not complete it. We removed two respondents because they had worked on over eighty apps while having less than three years of mobile development experience, one respondent who finished the survey in less than three minutes, and one respondent who did not pass the attention check question. In total, we received 328 valid responses from Prolific.

GitHub. We sent emails to GitHub users who contributed to the top 1,000 GitHub repositories (sorted by the number of stars) written either in (1) Java (with "Android" as an additional keyword), or (2) Objective-C or Swift (with "iOS" as an additional keyword). In total, we sent out 33,675 emails, out of which 128 started the survey, 51 respondents did not finish the survey, and five had not developed apps in the past three years. Other checks did not result in removing any additional responses. In total, we received 72 valid responses from GitHub emails. These participants were offered to provide an email to enter into a raffle for a £30 gift card for each 20 participants; 57 participants decided to enter the raffle, out of which three random participants received a gift card.

Other Channels. We made an effort to recruit women and minority groups by posting the survey in 20 LinkedIn groups specific to these populations. 14 respondents started the survey, seven did not finish the survey, and the other seven had not worked on any apps in the past three years. Therefore, we did not receive any valid responses from these channels.

The anonymized dataset for multiple-choice responses, excluding the open-ended responses (per participant consent), for the 400 valid participants is available online at DOI: 10.7488/ds/3045.

3.3 Data Analysis

3.3.1 Quantitative Analysis

We fitted a generalized linear mixed model with the binary value of choice between personalized (coded as 0) and non-personalized ads (coded as 1) as the dependent variable because each participant contributed two output values, one per app category. The model consisted of the six conditions (with Control as the baseline), app category (with gaming as the baseline), and several demographics as fixed effects, and participants as random effects, given that we had two data points per participant (gaming and financial apps) [73]. The regression analysis was conducted in R using the lme4 (`glmer`) [21] and arm [38] packages using binomial family (`logit` was the link function).

3.3.2 Qualitative Analysis

The count of words in the three open-ended questions showed that the answers were brief (on average 20 words, $SD = 16$) and enabled us to use affinity diagrams—a tool for organizing and consolidating output from a brainstorming session according to its affinity, or similarity—for analysis [24, 55]. We used the virtual collaboration platform Miro [72] to create separate boards for each open-ended question and posted virtual sticky notes with participants' responses. During a half-day virtual session with five security and privacy researchers with a minimum Master's degree in computer science, and one senior Android developer, we identified the common themes through group affinity diagram building.

3.4 Limitations

As with any self-reported data, respondents' survey answers may be subject to social desirability bias [36] and may differ from actual behaviors (so called, privacy paradox [54]). However, our use of role-playing scenarios and questions about intentions (rather than only attitudes) partially mitigates these biases, as intentions are shown to significantly correlate with behaviors [8, 32]. Our work complements and extends other privacy-related studies with developers [59, 101, 106] by conducting a controlled study with high internal validity which provides a foundation for future validation work. The results show a promising effect which will need further field experiments to fully test the generalizability.

Compared to other studies using similar recruitment strategies, the response rate for GitHub emails in our study is 0.21%, which is similar to 0.31% in [105] and lower than 1.3% in [1]. However, we were able to recruit a sufficient number of participants through Prolific. Moreover, mentioning ad networks in the recruitment email could deter people concerned about user privacy or ad networks. However, our results do not support that worry, demonstrating a wide variety of opinions about ad networks and user privacy.

Due to the demographic composition of the Prolific participant pool [35], our sample is predominately European, which could result in participants being more aware of European privacy laws, i.e. GDPR. However, GDPR's jurisdiction applies worldwide and many developers create apps for different geographic markets, mitigating this concern. To geographically balance our sample, we used additional Prolific screening criteria to exclude European countries for 274 respondents of the screening survey. The diverse geographic background of GitHub participants also added diversity to our sample. While our results may not be generalizable to all populations, it provides insights on the impact of various nudges on developers' decisions. Additionally, including geographic variable did not improve our model's fitness and did not reveal any significant relation to the outcome variable. Future research is encouraged to validate the results with other populations.

Identification of participants as developers was self-reported, as we did not test them. However, we believe it does not undermine the validity of results, as GitHub is a platform targeted at developers, and Prolific participants had previously marked themselves as having computer programming experience. The recruitment materials also highlighted that the study was about improving advertising library integration experience; such jargon is likely to defer participants without relevant experience and attract developers.

4 Results

We first report participants' demographics in Section 4.1, then the main experimental effects in Sections 4.2 and 4.3, and finally the additional findings about participants' opinions and attitudes about ads personalization in Section 4.4 to contextualize and interpret the main results.

4.1 Participants

Our participants are mostly European (66%), male (82%)—representative of the male-dominated profession [98], have on average 5.1 years of experience in software development ($SD = 5.3$), 2.7 years of experience in mobile development ($SD = 2.6$), on average worked on 3.5 apps in the past three years ($SD = 4.2$), 73% worked in software teams (e.g., developer, tester, or manager), 46% hold a software development position, 69% had previously integrated an ad library, and 78% make money from software development (see Table 5 in the Appendix). Over 90% of Google Play developers have one to nine apps under their account (as of 2015) [115], suggesting that our sample represents a portion of mobile developers. More than half (57%) of participants have used at least one ad network in their apps. Google AdMob (48%), Facebook Audience Network (20%), and Unity Ads (20%) were the most popular ad networks.
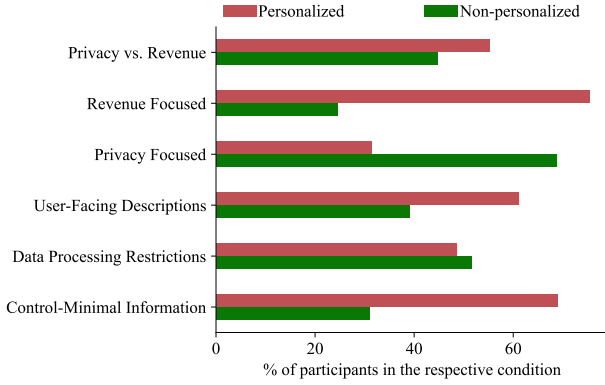
Figure 1: Participants' choices between personalized and non-personalized ads across the six conditions.

4.2 Choices Between Personalized and Non-Personalized Ads

As shown in Figure 1 (RQ1), the majority of participants chose personalized ads in the Revenue Focused (75%), Control (69%), and User-Facing Description (61%) conditions, and non-personalized ads in the Privacy Focused condition (69%). In the Data Processing Restrictions and Privacy vs. Revenue conditions, the choices between the two types of ads were split almost equally, with 49% and 55% respectively choosing the personalized ads.

The regression analysis (Table 1) confirms that the choice framing does impact participants' choices (RQ1). The strongest effect was in the Privacy Focused condition: using framing that explicitly mentions the implication for user privacy and what data will be used nudged participants to be 11.06 times ($p < .001$) more likely to choose non-personalized ads over personalized ads, compared to the Control condition. In the Data Processing Restrictions condition, framing that emphasized data restrictions associated with the choice of ads nudges participants to be 3.45 times ($p = .011$) more likely to choose the non-personalized ads compared to the Control condition. The results in the Revenue Focused, User-Facing Descriptions, and Privacy vs. Revenue conditions were not significantly different from the Control condition. In other words, using the neutral framing about personalized and non-personalized ads (Control condition), emphasizing the consequences of personalized ads on app's revenue (Revenue Focused condition), leveraging the user-facing description to provide transparency to users about whether app uses personalized ads based on users' personal data or not (User-Facing Description condition), and providing an explicit choice between user privacy and app's revenue (Privacy vs. Revenue) similarly affect participants' choices to integrate predominantly personalized ads in the apps.

Impact of App Category: Financial vs Gaming. Participants' choices between the app categories were not differ-

Independent Variables	ORs	CI (95%)	*p*–value
Condition			
Control–Minimal Information		Reference	
Data Processing Restrictions	3.45	1.32–8.98	.011*
User-Facing Descriptions	1.38	0.54–3.50	.502
Privacy Focused	11.06	3.97–30.75	<.001***
Revenue Focused	0.50	0.19–1.33	.164
Privacy vs. Revenue	2.48	0.97–6.35	.058
App Category			
Gaming app		Reference	
Financial app	1.02	0.70–1.49	.923
Given Priority to Privacy in Development Routines			
Low priority		Reference	
Not a priority	1.27	0.11–15.04	.851
Medium priority	1.84	0.75–4.51	.184
High priority	3.94	1.59–9.75	.003**
Essential	10.33	3.43–31.11	<.001***
Main Income Source			
Salary, not dependent on app revenue		Reference	
Don't make money from app development	2.63	1.23–5.66	.013*
Salary, partially dependent on app revenue	0.57	0.27–1.17	.126
Direct app revenue	0.73	0.32–1.66	.447
Other	0.90	0.07–11.17	.934
Years of experience in software development	1.08	1.02–1.14	.007**
Number of developed apps in the past three years	0.92	0.86–0.99	.033*
(Intercept)	0.09	0.03–0.3	< .001***

Table 1: Generalized linear mixed model regression. Outcome variable is the binary choice between personalized (coded as 0) and non-personalized ads (coded as 1). OR: odds ratios, CI: confidence intervals, conditional R^2: .614 (represents how much of the variance is explained by the model [62]), No. observations: 800, $*p < .05, **p < .01, ***p < .001$.

ent; 57% of participants chose personalized ads in both categories. Thus, our expectation that the financial app would trigger more privacy-preserving choices (non-personalized ads) because it carries obvious privacy risks for users is not supported by the data. We did not observe a significant interaction between conditions and app categories. In Section 5.3, we explore the potential reasons behind this effect based on participants' open-ended answers.

Impact of Demographics. We also included the demographic variables in the model that improved the model's fit. We found that participants, who consider privacy an essential or high priority are 10.33 ($p < .001$) and 3.94 times ($p = .003$), respectively, more likely to choose non-personalized ads compared to those who consider privacy a low priority in daily development routines (we selected the low priority as the reference category here because the not a priority category only had five responses making the category sizes highly unbalanced). Participants, who do not make money from software or apps, are 2.63 times ($p = .013$) more likely to choose the non-personalized ads compared to those

whose income is from software/app development but is not dependent on app revenue.

Each additional year of experience in software development increases the likelihood of choosing non-personalized ads by 8% ($p < .001$), but each additional app that participants developed in the past three years decreases the odds of choosing the non-personalized ads by 8% ($p = .033$).

The inverse relation between the number of developed apps and the choice of non-personalized ads may be related to the participants getting used to the status quo in that area as they develop more apps. More years of experience may also increase developers' awareness about other app monetisation methods. Inclusion of other variables, such as years of experience in mobile development, did not improve the model fit, thus we did not include them in the final model.

4.3 Reasons Behind the Ad Type Choices

Using affinity diagrams, as discussed in Section 3.3, we constructed themes around participants' responses to the question: "What was the biggest reason that made you pick the ad type: [their choice]" (RQ2). Table 2 shows the resulting themes. We provide the unique count of participants that mention each theme at all (out of 400) as well as the number of responses that mention a theme (out of 800) as each participant provided a response for each of the two apps. Quotes are labeled with P or NP based on the participant's choice for personalized or non-personalized ads. Theme frequencies are provided to give a sense of scale, but should not be used for generalization or statistical analysis since they only measure what participants thought to mention.

We identified three major reasons for choosing personalized or non-personalized ads: expected impact on revenue, user privacy, and relevance to users. Participants in the Privacy Focused condition mentioned privacy most often, and participants in the Revenue Focused condition mentioned monetisation most often as a reason for their ads choices.

Impact on Revenue. A main reason for choosing a certain ad type was related to monetisation goals and impact on revenue, mentioned by 41.5% of participants (166/400). Those, who chose personalized ads, were especially likely to relate their choice to expected positive impact on revenue (232/800): "To ensure most people click on the ad, increasing the apps revenue" (P309). Less often participants chose non-personalized ads with the expectations of positive impact on revenue (24/800): "I believe that providing non-customized ads would help to increase consumption regardless of the type of ad" (NP68).

User Privacy. Out of participants who chose non-personalized ads, most did it because of user privacy (269/800), for example, to protect users' sensitive data

(35/800), gain their trust (40/800), comply with privacy regulations (13/800), or gain a competitive advantage (12/800): "App doesn't have personalized information about the user. Also, it is easier to comply with GDPR rules that way" (NP213), "Given Apple's latest privacy changes, users are more aware of apps that invade their privacy and as a result, could be less likely to download these apps" (NP224). Some mentioned the long-term benefits of user trust over the short-term gains from violating user privacy: "Users trust in protecting the privacy is the most valuable good for a developer (besides quality of content). Aiming at a one-hit-wonder one wouldn't care about it, but with long time plans this is the only manageable compromise for all stakeholders" (NP135).

Participants, who mentioned privacy in relation to their choice of personalized ads (24/800), mostly assumed that users do not care about privacy (7/800): "Just like it is with facebook and other big ad circulators, It's proven that people only care about their privacy on a surface level" (P202).

Several participants acknowledged the trade-off between user privacy, trust, and other considerations such as revenue (6/400): "I was torn. On the one hand, personalized ads in the context of ones [sic] finances are going to have a *much* higher CPM and I would like to capitalize on that. However, because I'm running an app whose data is sensitive and where I am more dependent on long term trust from my users, I decided to make the ads less personalized to start so that I can have fewer scary disclosures and consent screens. If the app is successful, I can always explore personalizing them later" (NP197). Participants also expressed struggling with the trade-off between revenue and user privacy: "Desire to protect customers privacy. This was a tough one and I waffled back and forth. If it offered higher payout I would have selected this option" (NP317).

Only seven participants mentioned the potential security risks associated with personalized ads: "This type of app wants to give the user a sense of security so personalized ads might put someone off from using this app to manage their finances" (NP473).

Relevance to Users. Many participants believe that ads should be interesting, relevant, engaging, and useful to the users (156/400). On the one hand, they believe that such ads are beneficial to the users: "Personalized ads are appealing to the user, a person interested in a specific topic would rather see/read more about it than a random ad" (P169). Given that personalized ads are targeted to users' potential interests, most participants driven by that reason selected the personalized ads over non-personalized ones (197/800 vs. 29/800). A smaller group of participants chose non-personalized ads because they considered them relevant to users: "Using non-personalized ads, you have the luxury of inserting different ads of which some may get the attention of the users further increasing the interaction" (NP163). Some participants were even worried that relevant ads may distract users' attention

Theme	Condition (participants, $N = 400$)							Ad Type Choices (occurrences, $N = 800$)		
	Control	Data Processing Restrictions	User-Facing Descriptions	Privacy Focused	Revenue Focused	Privacy vs. Revenue	Total	Personalized	Non-Personalized	Total
Impact on revenue	32 (8.0%)	16 (4.0%)	29 (7.2%)	18 (4.5%)	46 (11.5%)	25 (6.2%)	166 (41.5%)	232 (29.0%)	24 (3.0%)	256 (32.0%)
User privacy	13 (3.2%)	34 (8.5%)	23 (5.8%)	48 (12.0%)	11 (2.8%)	32 (8.0%)	161 (40.2%)	24 (3.0%)	269 (33.6%)	293 (36.6%)
Sensitive data	1 (0.2%)	9 (2.2%)	4 (1.0%)	11 (2.8%)	1 (0.2%)	6 (1.5%)	32 (8.0%)	-	35 (4.4%)	35 (4.4%)
User trust	2 (0.5%)	6 (1.5%)	5 (1.2%)	3 (0.8%)	7 (1.8%)	7 (1.8%)	30 (7.5%)	5 (0.6%)	40 (5.0%)	45 (5.6%)
Compliance	-	6 (1.5%)	1 (0.2%)	2 (0.5%)	2 (0.5%)	1 (0.2%)	12 (3.0%)	3 (0.4%)	13 (1.6%)	16 (2.0%)
Competitive advantage	-	3 (0.8%)	-	3 (0.8%)	-	4 (1.0%)	10 (2.5%)	-	12 (1.5%)	12 (1.5%)
Users don't care about privacy	-	-	-	6 (1.5%)	-	1 (0.2%)	7 (1.8%)	7 (0.9%)	-	7 (0.9%)
Security reasons	1 (0.2%)	2 (0.5%)	1 (0.2%)	2 (0.5%)	-	1 (0.2%)	7 (1.8%)	1 (0.1%)	8 (1.0%)	9 (1.1%)
Privacy & ethics trade-off	-	-	1 (0.2%)	3 (0.8%)	1 (0.2%)	1 (0.2%)	6 (1.5%)	4 (0.5%)	4 (0.5%)	8 (1.0%)
Relevance to users	33 (8.2%)	26 (6.5%)	33 (8.2%)	11 (2.8%)	30 (7.5%)	23 (5.8%)	156 (39.0%)	197 (24.6%)	29 (3.6%)	226 (28.2%)
User experience	8 (2.0%)	9 (2.2%)	17 (4.2%)	12 (3.0%)	11 (2.8%)	3 (0.8%)	60 (15.0%)	48 (6.0%)	27 (3.4%)	75 (9.4%)
Category-related reasons	7 (1.8%)	9 (2.2%)	6 (1.5%)	18 (4.5%)	5 (1.2%)	15 (3.8%)	60 (15.0%)	18 (2.2%)	89 (11.1%)	107 (13.4%)
Finance-related	3 (0.8%)	9 (2.2%)	4 (1.0%)	13 (3.2%)	5 (1.2%)	8 (2.0%)	42 (10.5%)	7 (0.9%)	72 (9.0%)	79 (9.9%)
Gaming-related	3 (0.8%)	2 (0.5%)	3 (0.8%)	7 (1.8%)	-	8 (2.0%)	23 (5.8%)	10 (1.2%)	20 (2.5%)	30 (3.8%)
Specificity of a target audience	2 (0.5%)	-	1 (0.2%)	5 (1.2%)	4 (1.0%)	5 (1.2%)	17 (4.2%)	10 (1.2%)	10 (1.2%)	20 (2.5%)
Users should decide	2 (0.5%)	2 (0.5%)	5 (1.2%)	4 (1.0%)	2 (0.5%)	2 (0.5%)	17 (4.2%)	18 (2.2%)	8 (1.0%)	26 (3.2%)
Easier to develop	-	3 (0.8%)	1 (0.2%)	-	4 (1.0%)	-	8 (2.0%)	1 (0.1%)	9 (1.1%)	10 (1.2%)
Everyone does it	-	1 (0.2%)	1 (0.2%)	3 (0.8%)	1 (0.2%)	1 (0.2%)	7 (1.8%)	7 (0.9%)	-	7 (0.9%)
Unclear responses	5 (1.2%)	4 (1.0%)	2 (0.5%)	4 (1.0%)	1 (0.2%)	3 (0.8%)	19 (4.8%)	15 (1.9%)	11 (1.4%)	26 (3.2%)

Table 2: Constructed themes from participants' answers about the primary reason for choosing the ad type.

away from the app, reducing engagement: "You would get distracted if you saw a product that you like, the user could easily close the app and search that product" (NP42).

Participants in the Privacy Focused condition were least likely to mention the relevance of ads to the users (11/400), but we did not observe much difference among the other conditions (23–33/400).

User Experience. Some participants (60/400) mentioned the impact of ads on user experience as a reason for their choice. In contrast to the theme about relevance of ads emphasizing their utility and benefits to the users, this theme emphasizes the emotional and experiential impact of ads.

Participants who chose personalized ads (48/800) thought that they are less annoying, more enjoyable, and of higher quality: "To avoid frustrating customers with irrelevant to their interests ads that they will be forced to watch throw [sic] to play the game for free personalized ads are a great choice to make fun the rewarded video ad format" (P493), ". . . I would like the ads to feel native to the app so it is a more professional experience for the user and as such high quality and personalized ads would fit better for such an app" (P333). Participants who chose non-personalized ads (27/800) believed them to be less invasive and creepy: "I feel that personalized ads are too intrusive and creepy, so I would rather opt for non-personalized ads. . . . I don't want to scare away users" (NP330). Some participants preferred to reduce the number of ads in general to minimise the interruption of the main interaction with the app, especially in the gaming context: "Gaming isn't a prime state to be in to think about purchases. As someone with experience, ads feel like a break

in action in games and I would say its not worth the extra money overall" (NP396).

Category-Related. Some participants said their choice of ad type partially depends on the app category, the data it collects, or the specific user audience it targets (60/400). For instance, we already discussed earlier that perceived sensitivity of user data may raise privacy and trust concerns, especially in the context of a financial app, leading participants to choose non-personalized ads: "We're building a financial app after all. The data in there is sensitive and if there have to be ads, they should in no way track the user. Otherwise we'll loose trust faster than we can build the app" (NP136). Similarly, some participants thought that the data collected in the gaming app is not sensitive, justifying the use of personalized ads: "The information shared with a gaming type of application may be not as important to the consumer" (P301). Others thought that the data collected in the gaming app does not reveal personal information, and thus cannot be used for targeting, leading to the choice of non-personalized ads: "A Gaming app should not have any access to personal data, so personalized advertising is just not possible" (NP192).

On the other hand, a few participants (6/400) thought that the target audience of a financial app is particularly valuable to advertisers, due to their higher buying power, thus, promising a particularly high return on personalized advertising: "The target market for the app is an older and more affluent audience, therefore it is worth exploring to show the personalized ads to yield a higher revenue" (P474).

Other Themes. These themes were mentioned by a few participants, but still provide interesting insights. For instance, 17 participants said that they prefer to let *users* decide what types of ads they want to see. For example, participant P39 shifted the responsibility to users assuming that they know what information was used for customizing the ad, what the privacy implications are of such targeting, and what the appropriate tools are for controlling online tracking: "Because I bet on the smart mind of my client, he/she should know how ads work and should know whether if the ad is shown after seeing custom profiling data or not and to offer the choice to get tracked or not" (P39). Participant NP299 acknowledged that there is currently little transparency about the data practices in app stores, and that users may not pay attention to the disclosures with poor usability: "Somehow in google play they do not give at least warnings and most users install without first reading labels. The case is to leave that label so that the user reads or does not read it is aware of the type of advertising that is included with the application" (NP299).

Eight participants expected that it will be easier and faster to implement non-personalized ads: "Helps to get app on stores, we are not collecting personal information and it helps to pass faster" (NP12). Seven participants chose personalized ads simply because it is common and it is the status quo in app advertising: "Many of the apps that I use have this type of ad" (P484).

4.4 Opinions About Ad Networks, Privacy Regulations, and Consent

In this section we report the results from the exit survey that helped us further contextualize and interpret the main treatment effects, as later discussed in Section 5.

Perceived Control Over Ads. While the choices about ad networks' and apps' business models are often made by upper-level and middle management (Figure 4 in the Appendix), our participants feel involved in that decision-making process. Many participants have been involved at least a moderate amount in choosing ad networks (36%), configuring ads (46.7%), and integrating the code to enable in-app ads (47.5%) (Figure 5 in the Appendix). However, despite the involvement in selecting ad networks, participants mostly agree that developers have moderate (40.25%) or very little (32.75%) control over the data collection by those networks (Figure 6 in the Appendix); and end-users have even less control (Wilcoxon signed-rank test of perceived end-user control relative to developer control: $U = 8409$, $p < .001$).

Reasons for Not Including an Ad Network. More than half (69%) of participants have used at least one ad network in their apps. We asked the remaining 123 participants to explain why they did not include any ad networks in their

Reason for Not Including Ad Networks	#Participants
No need to monetize the app	50 (40.65%)
Generic reasons	31 (25.2%)
Paid apps	12 (9.8%)
Open-source or free apps	7 (5.7%)
Apps not intended for public audience	25 (20.3%)
Small and personal projects	17 (13.8%)
Academic projects	8 (6.5%)
Expected negative impact on user experience	18 (14.6%)
Decision was made by others	16 (13.0%)
It's a responsibility of others	7 (5.7%)
Don't know how to do it	5 (4.1%)
User privacy	4 (3.3%)
Still in early development stages	4 (3.3%)
Unclear responses	4 (3.3%)

Table 3: Constructed themes around participants' reasons for not including ad networks in their apps ($N = 123$).

apps and constructed themes around participants' answers (Table 3), as discussed in Section 3.

Forty percent of these participants (50/123) did not integrate ad networks because there was no need to use ads to monetize the app, for instance, because it was free or open-source, or relied on other sources of revenue. About 20% of participants (25/123) did not aim for a broad audience and public use, but used instead for small personal projects, learning experience, homework, or academic research. Some participants (18/123) considered ads intrusive and damaging to user experience: "I've always found it less intrusive for the end-users and a much smoother experience for them overall so buying a premium version would be preferred as a way to monetize the apps" (P131). Others (16/123) said that they did not have control over that decision, e.g., because they were developing an app for a client. A few participants said that they did not know how to integrate an ad network (5/123), it was someone else's responsibility to do it (7/123), or the project was still in the early development stage for ad integration (4/123). Only four participants explicitly mentioned concerns about user privacy: "Ad networks are not transparent and can't be audited. I can't control the amount of information fetch from my users" (P201).

Perceived Impact of Personalized Ads on Revenue and User Base. We asked participants how choosing personalized ads over non-personalized ads is likely to affect the revenue and number of users (Figure 7). The majority of participants expected an increase in revenue in both app categories, but no or little decrease in the user base. Specifically, almost half of participants expected an increase in revenue by

up to 40%. Slightly more participants believed that the user base won't change in the gaming app compared to financial app (43% vs. 32.5%). However, 16-18% of participants believed that deploying personalized ads will not change the revenue at all, or even *decrease* the revenue in both app categories, and decrease the user base by up to 40% in financial (32%) and gaming (23%) apps.

Beliefs About Privacy Regulations. In the survey scenarios, we told participants that the apps will be published in Europe and the United States and are mainly targeted towards adults above age of 18. For both apps, we asked participants to select the regulations that would apply to each app, providing both full names and abbreviations of all regulation options. Most participants (70.5%) correctly chose GDPR, while the American privacy regulation CCPA was not chosen as often (26%), although the app descriptions explicitly mentioned that the apps will be published in both European and American markets. Moreover, specialized American regulations—Children's Online Privacy Protection Act (COPPA) [28] and Health Insurance Portability and Accountability Act (HIPAA) [49]—were chosen by 22.8% and 9.9%, respectively, although the described apps were not directed at children and did not collect health-related information.

It is possible that the participants, most of which are from Europe, are more familiar with the European regulations than the American ones, however, we did not find a significant difference between the answers about applicable regulations between the European and North American residents (Mann-Whitney test: $U = 98708.0, p = 0.174$). Finally, 22.8% of participants did not know what regulations apply to the apps, and 2.9% thought that none of them apply. These results show that developers may not be familiar with privacy regulations outside their home country and may not know which regulations are applicable to their apps. It also echos the findings of interviews with developers that they rarely know about privacy guidelines and required measures for privacy [14].

Opinions About User Consent. In the exit survey, we asked participants how they would ask for user consent, assuming they had decided to use personalized ads (Table 5 in the Appendix). The majority (32%) selected the consent form provided by our imaginary Acme ad network. Others preferred to rely on the consent forms provided by leading tech companies (22.5%), such as Facebook or Google, or not-for-profit organizations (10.7%), such as Mozilla or Electronic Frontier Foundation, or use their own consent forms (17.7%). Only 9.75% said they will not ask for user consent at all, assuming that ad network or someone else in the team will take care of it, or because they find the process difficult, unfamiliar, unimportant, or simply not required. Finally, 6% said they would consult the specialized companies providing compliance services.

Information Source	#Participants
Reuse available materials	21 (29.6%)
From other companies and not-for-profits	17 (23.9%)
Ready-to-use templates	4 (5.6%)
Guidelines	14 (19.7%)
Legal policies (e.g., GDPR)	10 (14.1%)
UX guidelines	4 (5.6%)
Online search	9 (12.7%)
Legal teams	7 (9.9%)
Relying on own knowledge	6 (8.5%)
Don't know	6 (8.5%)
Unclear responses	12 (16.9%)

Table 4: Constructed themes around participant's information sources for building their consent forms ($N = 71$).

We asked the 71 participants, who indicated they would use their own consent form, what *information sources* they would use to build it (Table 4). After constructing themes around open-ended responses using affinity diagrams, we found that almost a third (29.6%) of participants would still fall back on the existing consent forms built by other teams, apps, companies, non-for-profit organizations, or ready-to-use templates, when building their own forms. Another 19.7% would use general guidelines, such as regulatory policies and recommendations; four participants mentioned using user experience guidelines and best practices when building consent forms: "Existing UX research on consent forms and how to maximize consent with storytelling" (P224).

Other participants said they would search for information about consent forms on the Internet (12.7%), rely on the legal teams or lawyers (9.9%), and their own knowledge or "common sense" (8.5%). However, what constitutes "common sense" for the developer may not necessarily represent what is "common sense" for users. For instance, P277 said that they would tell users that their app uses ads, but would refrain from disclosing that those ads are based on personal information about them: "I'd be upfront about including ads but not state that they dig into people's history" (P277). Finally, 8.5% said they do not know what information they would rely on when building consent forms.

5 Discussion and Future Work

Prior work suggests the importance of improving usability of *security*-related interfaces for developers, for example, through security APIs [43], security notifications [105], and providing secure code examples [69, 70, 71]. Our study highlights the importance of *privacy* interfaces as well by looking at the impact of choice framing on developers' decisions

about user privacy while interacting with ad networks. We hypothesize that the low rate of GDPR-compliant consent forms on websites [37, 67, 112] and the abundance of non-compliant Android apps [60, 89, 96, 118] may partially be caused by developers' low awareness about or consideration of consequences of their decisions on user privacy. We find that incorporating nudges in the design of developers' tools may assist developers in making decisions that consider user privacy in their software development processes.

5.1 Provide Information About Privacy Implications of Ad Personalization

The choice framing that described data processing as being restricted to contextual information instead of past behaviors produced positive but weaker effects compared to the explicit use of privacy labels (11.06 vs. 3.45 times increase in the likelihood to choose non-personalized ads). We believe that this is because in the former case participants had to evaluate themselves the implications of using contextual vs. behavioral targeting on user privacy, while labels that clearly indicated the positive and negative privacy consequences simplified this task. We hypothesize that developers may not fully understand the differences between contextual and behavioral targeting and associated privacy implications; future work is called to explore this hypothesis.

Thus, we recommend ad networks to include information to help developers evaluate privacy implications of their decisions in a transparent, concise, and direct way, by including clear privacy labels to the choices about the ad types. Including these options in the documentation and quick start guides as part of developers' workflow for ads integration may also assist developers in considering user privacy as part of their app development procedure. Additional information on users' concerns about behavioral targeting (e.g., discomforting [63, 117], discriminating [86], and intrusive [83, 88]) might facilitate developers' assessment of privacy implications or support the claims about their relative privacy invasiveness; future work is needed to study how to effectively integrate this information without making the choice text options longer, and whether the manipulation is effective in nudging developers' choices in a less controlled setting.

5.2 Improve the Effectiveness of User-Facing Privacy Descriptions

Prior work recommends emphasizing privacy features in the app stores [59], for instance, the recent inclusion of "Privacy Details" in the Apple App Store aimed at explaining apps' privacy practices before users download them [12]. However, our experiment did not find evidence that adding user-facing descriptions (with our choice framing) of app's ad targeting practices would nudge participants to integrate less invasive non-personalized ads. Participants' open-ended comments

suggest a potential explanation: most participants do not expect personalized ads to reduce their app's user base; they also believe that personalized ads are more relevant and less annoying to the users. In other words, some participants believed that telling users that an app shows ads tailored to their personal information will not discourage users from downloading it, and indeed, may even attract users who prefer ads relevant and customized to their interests. However, prior work shows that some users do not like behaviorally targeted ads, find them invasive and creepy, and try to avoid or block such ads [5, 10, 75, 97, 99, 110].

Future work is called for to explore more efficient ways to nudge developers to consider privacy implications of their in-app ad choices. For instance, studying how to best provide evidence to developers about user opinions around ads, privacy preferences, and the impact of app-store presented information, would all help better inform developers' choices. Moreover, future work may test and improve the effectiveness of the existing ways to increase transparency and developers' responsibility to users' regarding their privacy, such as adding "Privacy Details" in the Apple App Store [12], potentially from a privacy nutrition labels perspective [53].

5.3 Reconcile Contradicting Beliefs

As we explained in Section 4.2, the app category did not impact the decisions between the personalized and non-personalized ads, and the number of participants in each group differed only slightly. The analysis of category-related reasons (Section 4.3) provides a potential explanation why we might have not observed a difference. Specifically, it revealed the contradicting beliefs about the same app category that lead to different ad type choices, potentially canceling out the effects of app category. For example, while some participants preferred non-personalized ads for financial apps to avoid raising privacy and trust concerns among users, others preferred to maximize profit from showing the personalized ads to this affluent user group, particularly valued by the advertisers. In the gaming context, because presumably the app does not collect sensitive information, some chose personalized ads as they believed it would not raise privacy concerns, others chose non-personalized ads as it would not be possible to customize ads due to the lack of personal information.

Similar contradictions are revealed in the experimental conditions. When we emphasized privacy implications, the majority of participants chose more privacy-friendly non-personalized ads. When we emphasized the implications on app's revenue, the majority chose revenue-maximizing personalized ads. However, when faced with an explicit choice between user privacy and app's revenue, the choices between two types of ads split almost equally, with a small preference for non-personalized ads. This finding suggests the balance between the contradicting values is fragile and can be easily manipulated. Similar to users' privacy decisions being

context-dependent [2, 4, 80], developers' decisions may also be driven by contextual factors. As some of our participants clarified in the open-ended responses, this choice may change depending on the associated impact on revenue or user privacy. For instance, if the promised increase in revenue is high enough, developers may choose it over user privacy; if they believe that the data collected by the app or context of the app in general is particularly sensitive to raise user concerns, they may be more prone to choose user privacy over profit.

Developers may integrate ad networks primarily because they see it as the only feasible way to monetize the app [68]. The current choice framing in the ad networks also favors the revenue and uses a language that nudges developers into choosing the personalized ads [102]. However, there are also hidden costs of mobile ads that many developers do not consider in weighing the trade-offs, such as frequent updating of ad-related code, and increased consumption of energy and network data on users' phone and subsequent decrease in app's use [44]. Future work could suggest ways to provide transparency about such trade-offs by looking at proposed frameworks for improving the equilibrium between the revenue and user privacy in smartphones by adjusting the level of privacy protection in response to ad-generated revenue [57].

Our results also inform regulators that slight changes in ad networks' interface design for developers may affect the fragile balance between the contradicting values of personalized ads and significantly affect developers' choices to benefit platform's interests in profit maximization. We recommend regulators build clear technical recommendations for providing choices to users, and to enforce that ad networks and other platforms use the mandated framing to promote users' welfare, and avoid effects driven by platforms' sole interests. Future work could provide inputs to the regulators by studying the usability of developer-facing interfaces (e.g., the privacy dashboard on Google AdMob), to inform the design of such interfaces and to provide suggestions to regulators on how to minimize the use of dark patterns in these interfaces.

5.4 Increase Developers' and Users' Control Over Data and Transparency

Many participants said that they do not have full control over ad networks' data collection and processing for ad personalization, and that users have even less control over it. We recommend ad networks, and app stores in particular, to increase the transparency about data practices, accountability to users, and developers' and users' control over data. For instance, Google Play's privacy nudges for permissions has shown success in reducing the number of permissions that developers request [85]. This model might be used to make information about third-party libraries such as ad networks more specific. We suggest app stores to scan for ad libraries and inform developers about their privacy implications during the automatic reviews of the apps (as they currently do for

other purposes such as displaying third-party apps [13]).

Some of our participants said that they prefer to let users decide what types of ads they want to see (personalized or non-personalized). However, this line of thought is not completely fair to the users in the environment of information asymmetry, where users are poorly informed about the data practices of apps and ad networks, and personal data flows are not transparent to the users [9, 15, 22, 78]. Thus, providing means for users to see what ad networks are being used in apps when installing a new app [29], what types of ads do the apps serve, and what personal information is used to customize them, as well as other improvement in user interfaces described in Section 5.2, might be effective. Prior results from user research may also help build usable privacy interfaces for developers and increase transparency and control. For instance, several elements of the labels such as data collection, purpose, and data sharing [34, 53] might be reused to inform developers about an ad network's data collection. Other proposed interfaces that visually represent permissions, purposes, data leaks [61, 114], data flows, the effects of removing and adding libraries [113], and integrating privacy checks into programming interfaces [58] might further inform developers about the privacy consequences of their choices. Not-for-profit organizations could build open-source services and easy-to-integrate privacy consent mechanisms to facilitate consent integration, and offer alternatives to for-profit large companies consent forms. Future work could also evaluate the effectiveness of various types of information sources on developers' success in building compliant and user-friendly consent forms (Table 4).

6 Conclusion

We present the results of an online experiment with 400 participants with mobile app development experience on their decisions regarding configuring ads for hypothetical apps. We find that the choice framing in ad networks significantly impacts developers' choices and subsequently privacy of millions of users. Thus, more control and transparency should be provided to developers and users in choosing the type of ads and data collection practices. Moreover, some of our participants incorrectly identified what privacy regulations would apply to the apps, and many said they rely on ad networks and examples from tech companies, when building user consent forms. This means that those companies are not only responsible to their own users, but also set example for other smaller companies and independent developers, further illustrating the large impact of ad network platform's design and choice framing on data practices in app development. Our results have implications for ad networks, app stores, and regulators by giving them grounds for promoting user privacy by improving the usability of developer-facing interfaces to empower developers in making informed decisions for their users.

Acknowledgments

References

[1] Yasemin Acar, Christian Stransky, Dominik Wermke, Michelle L. Mazurek, and Sascha Fahl. "Security Developer Studies with GitHub Users: Exploring a Convenience Sample". In: *Thirteenth Symposium on Usable Privacy and Security (SOUPS 2017)*. Santa Clara, CA: USENIX Association, July 2017, pp. 81–95. URL: https://www.usenix.org/conference/soups2017/technical-sessions/presentation/acar.

[2] Mark Ackerman, Trevor Darrell, and Daniel J Weitzner. "Privacy in context". In: *Human–Computer Interaction* 16.2-4 (2001), pp. 167–176. DOI: 10.1207/S15327051HCI16234_03.

[3] Alessandro Acquisti, Idris Adjerid, Rebecca Balebako, Laura Brandimarte, Lorrie Faith Cranor, Saranga Komanduri, Pedro Giovanni Leon, Norman Sadeh, Florian Schaub, Manya Sleeper, Yang Wang, and Shomir Wilson. "Nudges for Privacy and Security: Understanding and Assisting Users' Choices Online". In: *ACM Computing Surveys* 50.3 (Aug. 2017). DOI: 10.1145/3054926.

[4] Alessandro Acquisti, Leslie K John, and George Loewenstein. "What is privacy worth?" In: *The Journal of Legal Studies* 42.2 (2013), pp. 249–274. DOI: 10.1086/671754.

[5] *Ad-Blocking: A deep-dive into ad-blocking trends*. Tech. rep. GlobalWebIndex, 2018. URL: https://www.globalwebindex.com/hubfs/Downloads/Ad-Blocking-trends-report.pdf (visited on 02/2021).

[6] Md Ahasanuzzaman, Safwat Hassan, Cor-Paul Bezemer, and Ahmed E. Hassan. "A longitudinal study of popular ad libraries in the Google Play Store". en. In: *Empirical Software Engineering* 25.1 (Jan. 2020), pp. 824–858. DOI: 10.1007/s10664-019-09766-x.

[7] Md Ahasanuzzaman, Safwat Hassan, and Ahmed E. Hassan. "Studying Ad Library Integration Strategies of Top Free-to-Download Apps". In: *IEEE Transactions on Software Engineering* PP (Mar. 2020), pp. 1–1. DOI: 10.1109/TSE.2020.2983399.

[8] Icek Ajzen. "From Intentions to Actions: A Theory of Planned Behavior". In: *Action Control: From Cognition to Behavior*. Ed. by Julius Kuhl and Jürgen Beckmann. Berlin, Heidelberg: Springer Berlin Heidelberg, 1985, pp. 11–39. DOI: 10.1007/978-3-642-69746-3_2.

[9] Urs-Vito Albrecht. "Transparency of health-apps for trust and decision making." eng. In: *Journal of medical Internet research* 15.12 (Dec. 2013). ISSN: 1438-8871 1439-4456. DOI: 10.2196/jmir.2981.

[10] Mimi An. *Why People Block Ads (And What It Means for Marketers and Advertisers)*. HubSpot. 2020. URL: https://blog.hubspot.com/marketing/why-people-block-ads-and-what-it-means-for-marketers-and-advertisers (visited on 02/2021).

[11] *Android Ad Network statistics and market share*. AppBrain. 2020. URL: https://www.appbrain.com/stats/libraries/ad-networks (visited on 09/2020).

[12] *App privacy details on the App Store*. Apple. 2021. URL: https://developer.apple.com/app-store/app-privacy-details/ (visited on 02/2021).

[13] *App Store Review Guidelines*. Apple. 2021. URL: https://developer.apple.com/app-store/review/guidelines/#unacceptable (visited on 02/2021).

[14] Rebecca Balebako and Lorrie Cranor. "Improving App Privacy: Nudging App Developers to Protect User Privacy". In: *IEEE Security Privacy* 12.4 (2014), pp. 55–58. DOI: 10.1109/MSP.2014.70.

[15] Rebecca Balebako, Jaeyeon Jung, Wei Lu, Lorrie Faith Cranor, and Carolyn Nguyen. ""Little Brothers Watching You": Raising Awareness of Data Leaks on Smartphones". In: *Proceedings of the Ninth Symposium on Usable Privacy and Security*. SOUPS '13. Newcastle, United Kingdom: ACM, 2013. DOI: 10.1145/2501604.2501616.

[16] Rebecca Balebako, Pedro G Leon, Hazim Almuhimedi, Patrick Gage Kelley, Jonathan Mugan, Alessandro Acquisti, Lorrie Cranor, and Norman Sadeh-Koniecpol. "Nudging users towards privacy on mobile devices". In: *CHI 2011 Workshop on Persuasion, Nudge, Influence and Coercion*. Carnegie Mellon University, 2011. DOI: 10.1184/R1/13028258.v1.

[17] Rebecca Balebako, Abigail Marsh, Jialiu Lin, Jason I Hong, and Lorrie Cranor. "The privacy and security behaviors of smartphone app developers". In: *Workshop on Usable Security (USEC'14)*. Internet Society, 2014. DOI: 10.14722/usec.2014.23006.

[18] Rebecca Balebako, Florian Schaub, Idris Adjerid, Alessandro Acquisti, and Lorrie Cranor. "The Impact of Timing on the Salience of Smartphone App Privacy Notices". In: *Proceedings of the 5th Annual ACM CCS Workshop on Security and Privacy in Smartphones and Mobile Devices*. SPSM '15. Denver, Colorado, USA: ACM, 2015, pp. 63–74. DOI: 10.1145/2808117.2808119.

[19] Kenneth A Bamberger, Serge Egelman, Catherine Han, Amit Elazari Bar On, and Irwin Reyes. "Can You Pay For Privacy? Consumer Expectations and the Behavior of Free and Paid Apps". In: *Berkeley Technology Law Journal* 35 (2020). DOI: 10.15779/Z38XP6V40J.

[20] Hyejin Bang and Bartosz W. Wojdynski. "Tracking users' visual attention and responses to personalized advertising based on task cognitive demand". In: *Computers in Human Behavior* 55 (2016), pp. 867–876. DOI: 10.1016/j.chb.2015.10.025.

[21] Douglas Bates, Martin Mächler, Ben Bolker, and Steve Walker. "Fitting Linear Mixed-Effects Models Using lme4". In: *Journal of Statistical Software, Articles* 67.1 (2015), pp. 1–48. DOI: 10.18637/jss.v067.i01.

[22] Jan Hendrik Betzing, Matthias Tietz, Jan vom Brocke, and Jörg Becker. "The impact of transparency on mobile privacy decision making". In: *Electronic Markets* 30.3 (Sept. 2020), pp. 607–625. ISSN: 1422-8890. DOI: 10.1007/s12525-019-00332-3.

[23] Christoph Bösch, Benjamin Erb, Frank Kargl, Henning Kopp, and Stefan Pfattheicher. "Tales from the Dark Side: Privacy Dark Strategies and Privacy Dark Patterns". In: *Proceedings on Privacy Enhancing Technologies* 2016.4 (2016), pp. 237–254. DOI: 10.1515/popets-2016-0038.

[24] Virginia Braun and Victoria Clarke. "Using thematic analysis in psychology". In: *Qualitative Research in Psychology* 3.2 (2006), pp. 77–101. DOI: 10.1191/1478088706qp063oa.

[25] Alex Braunstein, Laura Granka, and Jessica Staddon. "Indirect Content Privacy Surveys: Measuring Privacy without Asking about It". In: *Proceedings of the Seventh Symposium on Usable Privacy and Security*. SOUPS '11. Pittsburgh, Pennsylvania: ACM, 2011. DOI: 10.1145/2078827.2078847.

[26] *California Consumer Privacy Act (CCPA)*. State of California Department of Justice. 2018. URL: https://oag.ca.gov/privacy/ccpa (visited on 09/2020).

[27] Ana Caraban, Evangelos Karapanos, Daniel Gonçalves, and Pedro Campos. "23 Ways to Nudge: A Review of Technology-Mediated Nudging in Human-Computer Interaction". In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI '19. Glasgow, Scotland UK: ACM, 2019, pp. 1–15. DOI: 10.1145/3290605.3300733.

[28] *Children's Online Privacy Protection Rule (COPPA)*. Federal Trade Commission. 1998. URL: https://www.ftc.gov/enforcement/rules/rulemaking-regulatory-reform-proceedings/childrens-online-privacy-protection-rule (visited on 09/2020).

[29] Saksham Chitkara, Nishad Gothoskar, Suhas Harish, Jason I. Hong, and Yuvraj Agarwal. "Does This App Really Need My Location? Context-Aware Privacy Management for Smartphones". In: *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 1.3 (Sept. 2017). DOI: 10.1145/3132029.

[30] Joseph Cox. *How the U.S. Military Buys Location Data from Ordinary Apps*. VICE. 2020. URL: https://www.vice.com/en/article/jgqm5x/us-military-location-data-xmode-locate-x (visited on 02/2021).

[31] *Digital Advertising Alliance (DAA)*. Digital Advertising Alliance (DAA). 2021. URL: https://digitaladvertisingalliance.org (visited on 02/2021).

[32] Serge Egelman, Marian Harbach, and Eyal Peer. "Behavior Ever Follows Intention? A Validation of the Security Behavior Intentions Scale (SeBIS)". In: *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. New York, NY, USA: ACM, 2016, pp. 5257–5261. DOI: 10.1145/2858036.2858265.

[33] Anirudh Ekambaranathan, Jun Zhao, and Max Van Kleek. "Understanding Value and Design Choices Made by Android Family App Developers". In: *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*. CHI EA '20. Honolulu, HI, USA: ACM, 2020, pp. 1–10. DOI: 10.1145/3334480.3383064.

[34] Pardis Emami-Naeini, Yuvraj Agarwal, Lorrie Faith Cranor, and Hanan Hibshi. "Ask the Experts: What Should Be on an IoT Privacy and Security Label?" In: *2020 IEEE Symposium on Security and Privacy (SP)*.

2020, pp. 447–464. DOI: 10.1109/SP40000.2020. 00043.

[35] *Explore our participant pool demographics*. Prolific. 2021. URL: https://www.prolific.co/ demographics/ (visited on 02/2021).

[36] Robert J Fisher. "Social desirability bias and the validity of indirect questioning". In: *Journal of consumer research* 20.2 (1993), pp. 303–315. DOI: 10.1086/ 209351.

[37] Imane Fouad, Cristiana Santos, Feras Al Kassar, Nataliia Bielova, and Stefano Calzavara. "On Compliance of Cookie Purposes with the Purpose Specification Principle". In: *IWPE 2020 - International Workshop on Privacy Engineering*. Genova, Italy, Sept. 2020, pp. 1–8. URL: https://hal.inria.fr/hal- 02567022.

[38] Andrew Gelman, Masanao Yajima Yu-Sung Su, Jennifer Hill, Maria Grazia Pittau, Jouni Kerman, Tian Zheng, and Vincent Dorie. *arm: Data Analysis Using Regression and Multilevel/Hierarchical Models*. 2020. URL: https://cran.r-project.org/package= arm (visited on 02/2021).

[39] *General Data Protection Regulation (GDPR)*. The European parliament and the council of the European union. 2018. URL: https://eur-lex.europa. eu/legal-content/EN/TXT/PDF/?uri=CELEX: 32016R0679 (visited on 09/2020).

[40] Phillipa Gill, Vijay Erramilli, Augustin Chaintreau, Balachander Krishnamurthy, Konstantina Papagiannaki, and Pablo Rodriguez. "Follow the Money: Understanding Economics of Online Aggregation and Advertising". In: *Proceedings of the 2013 Conference on Internet Measurement Conference*. IMC '13. Barcelona, Spain: ACM, 2013, pp. 141–148. DOI: 10.1145/2504730.2504768.

[41] Avi Goldfarb. "What is Different About Online Advertising?" In: *Review of Industrial Organization* 44.2 (Mar. 2014), pp. 115–129. DOI: 10.1007/s11151- 013-9399-3.

[42] Colin M. Gray, Yubo Kou, Bryan Battles, Joseph Hoggatt, and Austin L. Toombs. "The Dark (Patterns) Side of UX Design". In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI '18. Montreal QC, Canada: ACM, 2018, pp. 1– 14. DOI: 10.1145/3173574.3174108.

[43] Matthew Green and Matthew Smith. "Developers Are Not the Enemy!: The Need for Usable Security APIs". In: *IEEE Security and Privacy* 14.5 (Sept. 2016), pp. 40–46. DOI: 10.1109/MSP.2016.111.

[44] Jiaping Gui, Stuart Mcilroy, Meiyappan Nagappan, and William G. J. Halfond. "Truth in Advertising: The Hidden Cost of Mobile Ads for Software Developers". In: *Proceedings of the 37th International Conference on Software Engineering - Volume 1*. ICSE '15. Florence, Italy: IEEE Press, 2015, pp. 100–110. DOI: 10.1109/ICSE.2015.32.

[45] Hana Habib, Sarah Pearman, Jiamin Wang, Yixin Zou, Alessandro Acquisti, Lorrie Faith Cranor, Norman Sadeh, and Florian Schaub. ""It's a Scavenger Hunt": Usability of Websites' Opt-Out and Data Deletion Choices". In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. CHI '20. Honolulu, HI, USA: ACM, 2020, pp. 1–12. DOI: 10.1145/3313831.3376511.

[46] Hana Habib, Yixin Zou, Aditi Jannu, Neha Sridhar, Chelse Swoopes, Alessandro Acquisti, Lorrie Faith Cranor, Norman Sadeh, and Florian Schaub. "An Empirical Analysis of Data Deletion and Opt-Out Choices on 150 Websites". In: *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*. Santa Clara, CA: USENIX Association, Aug. 2019. URL: https://www.usenix.org/conference/ soups2019/presentation/habib.

[47] Catherine Han, Irwin Reyes, Álvaro Feal, Joel Reardon, Primal Wijesekera, Amit Elazari, Kenneth A Bamberger, and Serge Egelman. "The Price is (Not) Right: Comparing Privacy in Free and Paid Apps". In: *Privacy Enhancing Technologies Symposium*. 2020, p. 21. DOI: 10.2478/popets-2020-0050.

[48] Boyuan He, Haitao Xu, Ling Jin, Guanyu Guo, Yan Chen, and Guangyao Weng. "An Investigation into Android In-App Ad Practice: Implications for App Developers". In: *IEEE INFOCOM 2018 - IEEE Conference on Computer Communications*. 2018, pp. 2465–2473. DOI: 10.1109/INFOCOM.2018. 8486010.

[49] *Health Insurance Portability and Accountability Act of 1996 (HIPAA)*. U.S. Department of Health & Human Services. 1996. URL: https://www.cdc.gov/ phlp/publications/topic/hipaa.html (visited on 02/2021).

[50] Maximilian Hils, Daniel W. Woods, and Rainer Böhme. "Measuring the Emergence of Consent Management on the Web". In: *Proceedings of the ACM Internet Measurement Conference*. IMC '20. Virtual Event, USA: ACM, 2020, pp. 317–332. DOI: 10. 1145/3419394.3423647.

[51] Ling Jin, Boyuan He, Guangyao Weng, Haitao Xu, Yan Chen, and Guanyu Guo. "MAdLens: Investigating into Android In-App Ad Practice at API Granularity". In: *IEEE Transactions on Mobile Computing*

PP.PP (2019), pp. 1–1. DOI: 10.1109/TMC.2019.2953609.

[52] Eric J Johnson, Steven Bellman, and Gerald L Lohse. "Defaults, Framing and Privacy: Why Opting In-Opting Out[1]". In: *Marketing letters* 13.1 (Feb. 2002), pp. 5–15. DOI: 10.1023/A:1015044207315.

[53] Patrick Gage Kelley, Joanna Bresee, Lorrie Faith Cranor, and Robert W. Reeder. "A "Nutrition Label" for Privacy". In: *Proceedings of the 5th Symposium on Usable Privacy and Security*. SOUPS '09. Mountain View, California, USA: ACM, 2009. DOI: 10.1145/1572532.1572538.

[54] Spyros Kokolakis. "Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon". In: *Computers & Security* 64 (2017), pp. 122–134. DOI: 10.1016/j.cose.2015.07.002.

[55] Jonathan Lazar, Jinjuan Heidi Feng, and Harry Hochheiser. "Chapter 8 - Interviews and focus groups". In: *Research Methods in Human Computer Interaction*. Ed. by Jonathan Lazar, Jinjuan Heidi Feng, and Harry Hochheiser. Second Edition. Boston: Morgan Kaufmann, 2017, pp. 187–228. DOI: 10.1016/B978-0-12-805390-4.00008-X.

[56] Pedro Leon, Blase Ur, Richard Shay, Yang Wang, Rebecca Balebako, and Lorrie Cranor. "Why Johnny Can't Opt out: A Usability Evaluation of Tools to Limit Online Behavioral Advertising". In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '12. Austin, Texas, USA: ACM, 2012, pp. 589–598. DOI: 10.1145/2207676.2207759.

[57] Ilias Leontiadis, Christos Efstratiou, Marco Picone, and Cecilia Mascolo. "Don't Kill My Ads! Balancing Privacy in an Ad-Supported Mobile Application Market". In: *Proceedings of the Twelfth Workshop on Mobile Computing Systems & Applications*. HotMobile '12. San Diego, California: ACM, 2012. DOI: 10.1145/2162081.2162084.

[58] Tianshi Li, Yuvraj Agarwal, and Jason I. Hong. "Coconut: An IDE Plugin for Developing Privacy-Friendly Apps". In: *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2.4 (Dec. 2018). DOI: 10.1145/3287056.

[59] Tianshi Li, Elizabeth Louie, Laura Dabbish, and Jason I. Hong. "How Developers Talk About Personal Data and What It Means for User Privacy: A Case Study of a Developer Forum on Reddit". In: *Proc. ACM Hum.-Comput. Interact.* 4.CSCW (Jan. 2021). DOI: 10.1145/3432919.

[60] Ilaria Liccardi, Monica Bulger, Hal Abelson, Daniel Weitzner, and Wendy Mackay. "Can apps play by the COPPA Rules?" In: *2014 Twelfth Annual International Conference on Privacy, Security and Trust*. 2014, pp. 1–9. DOI: 10.1109/PST.2014.6890917.

[61] Bin Liu, Mads Schaarup Andersen, Florian Schaub, Hazim Almuhimedi, Shikun (Aerin) Zhang, Norman Sadeh, Yuvraj Agarwal, and Alessandro Acquisti. "Follow My Recommendations: A Personalized Privacy Assistant for Mobile App Permissions". In: *Twelfth Symposium on Usable Privacy and Security (SOUPS 2016)*. Denver, CO: USENIX Association, June 2016, pp. 27–41. URL: https://www.usenix.org/conference/soups2016/technical-sessions/presentation/liu.

[62] Daniel Lüdecke, Dominique Makowski, Philip Waggoner, and Indrajeet Patil. *performance: Assessment of Regression Models Performance*. R package. 2020. DOI: 10.5281/zenodo.3952174.

[63] Miguel Malheiros, Charlene Jennett, Snehalee Patel, Sacha Brostoff, and Martina Angela Sasse. "Too Close for Comfort: A Study of the Effectiveness and Acceptability of Rich-Media Personalized Advertising". In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. CHI '12. Austin, Texas, USA: ACM, 2012, pp. 579–588. DOI: 10.1145/2207676.2207758.

[64] Miriam Marciel, J. G. Cabañas, Y. Kassa, R. Gonzalez, and M. Ahmed. "The Value of Online Users: Empirical Evaluation of the Price of Personalized Ads". In: *2016 11th International Conference on Availability, Reliability and Security (ARES)*. 2016, pp. 694–700. DOI: 10.1109/ARES.2016.89.

[65] Arunesh Mathur, Gunes Acar, Michael J. Friedman, Elena Lucherini, Jonathan Mayer, Marshini Chetty, and Arvind Narayanan. "Dark Patterns at Scale: Findings from a Crawl of 11K Shopping Websites". In: *Proceedings of the ACM on Human-Computer Interaction* 3.CSCW (Nov. 2019). DOI: 10.1145/3359183.

[66] Arunesh Mathur, Jessica Vitak, Arvind Narayanan, and Marshini Chetty. "Characterizing the Use of Browser-Based Blocking Extensions To Prevent Online Tracking". In: *Fourteenth Symposium on Usable Privacy and Security (SOUPS 2018)*. Baltimore, MD: USENIX Association, Aug. 2018, pp. 103–116. URL: https://www.usenix.org/conference/soups2018/presentation/mathur.

[67] Celestin Matte, Nataliia Bielova, and Cristiana Santos. "Do Cookie Banners Respect my Choice? : Measuring Legal Compliance of Banners from IAB Europe's Transparency and Consent Framework". In:

2020 IEEE Symposium on Security and Privacy (SP). May 2020, pp. 791–809. DOI: 10.1109/SP40000.2020.00076.

[68] Abraham H. Mhaidli, Yixin Zou, and Florian Schaub. ""We Can't Live Without Them!" App Developers' Adoption of Ad Networks and Their Considerations of Consumer Risks". In: *Fifteenth Symposium on Usable Privacy and Security (SOUPS 2019)*. Santa Clara, CA: USENIX Association, Aug. 2019. URL: https://www.usenix.org/conference/soups2019/presentation/mhaidli.

[69] Kai Mindermann, Philipp Keck, and Stefan Wagner. "How Usable Are Rust Cryptography APIs?" In: *2018 IEEE International Conference on Software Quality, Reliability and Security (QRS)*. IEEE, July 2018, pp. 143–154. DOI: 10.1109/qrs.2018.00028.

[70] Kai Mindermann and Stefan Wagner. "Fluid Intelligence Doesn't Matter! Effects of Code Examples on the Usability of Crypto APIs". In: *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering: Companion Proceedings*. ICSE '20. Seoul, South Korea: ACM, 2020, pp. 306–307. DOI: 10.1145/3377812.3390892.

[71] Kai Mindermann and Stefan Wagner. "Usability and Security Effects of Code Examples on Crypto APIs". In: *2018 16th Annual Conference on Privacy, Security and Trust (PST)*. IEEE, Aug. 2018, pp. 1–2. DOI: 10.1109/PST.2018.8514203.

[72] *Miro | Online Whiteboard for Visual Collaboration*. Miro. 2021. URL: https://miro.com/ (visited on 02/2021).

[73] *Mixed Effects Logistic Regression*. UCLA: Statistical Consulting Group. 2020. URL: https://stats.idre.ucla.edu/stata/dae/mixed-effects-logistic-regression/ (visited on 02/2021).

[74] *MoPub Integration Suite*. Twitter MoPub. 2021. URL: https://developers.mopub.com/publishers/integrate/ (visited on 02/2021).

[75] Lymari Morales. *U.S. Internet Users Ready to Limit Online Tracking for Ads*. Gallup Polls. 2010. URL: https://news.gallup.com/poll/145337/internet-users-ready-limit-online-tracking-ads.aspx (visited on 02/2021).

[76] *Most popular Apple App Store categories in August 2020, by share of available apps*. Statista. 2020. URL: https://www.statista.com/statistics/270291/popular-categories-in-the-app-store/ (visited on 02/2021).

[77] *Most popular Google Play app categories as of 3rd quarter 2020, by share of available apps*. Statista. 2020. URL: https://www.statista.com/statistics/279286/google-play-android-app-categories/ (visited on 02/2021).

[78] Patrick Murmann. "Usable Transparency for Enhancing Privacy in Mobile Health Apps". In: *Proceedings of the 20th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct*. MobileHCI '18. Barcelona, Spain: ACM, 2018, pp. 440–442. DOI: 10.1145/3236112.3236184.

[79] Arvind Narayanan, Arunesh Mathur, Marshini Chetty, and Mihir Kshirsagar. "Dark Patterns: Past, Present, and Future". In: *Queue* 18.2 (Apr. 2020), pp. 67–92. DOI: 10.1145/3400899.3400901.

[80] Helen Nissenbaum. "Privacy as contextual integrity". In: *Wash. L. Rev.* 79 (2004), p. 119.

[81] Midas Nouwens, Ilaria Liccardi, Michael Veale, David Karger, and Lalana Kagal. "Dark Patterns after the GDPR: Scraping Consent Pop-Ups and Demonstrating Their Influence". In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. CHI '20. Honolulu, HI, USA: ACM, 2020, pp. 1–13. DOI: 10.1145/3313831.3376321.

[82] *Number of software developers worldwide in 2018, 2019, 2023 and 2024*. Statista. 2020. URL: https://www.statista.com/statistics/627312/worldwide-developer-population/ (visited on 02/2021).

[83] Katie O'Donnell and Henriette Cramer. "People's Perceptions of Personalized Ads". In: *Proceedings of the 24th International Conference on World Wide Web*. WWW '15 Companion. Florence, Italy: ACM, 2015, pp. 1293–1298. DOI: 10.1145/2740908.2742003.

[84] *Out of Control - How consumers are exploited by the online advertising industry*. Forbrukerrådet. 2020. URL: https://www.forbrukerradet.no/side/new-study-the-advertising-industry-is-systematically-breaking-the-law (visited on 09/2020).

[85] Sai Teja Peddinti, Igor Bilogrevic, Nina Taft, Martin Pelikan, Úlfar Erlingsson, Pauline Anthonysamy, and Giles Hogben. "Reducing Permission Requests in Mobile Apps". In: *Proceedings of the Internet Measurement Conference*. IMC '19. Amsterdam, Netherlands: ACM, 2019, pp. 259–266. DOI: 10.1145/3355369.3355584.

[86] Angelisa C. Plane, Elissa M. Redmiles, Michelle L. Mazurek, and Michael Carl Tschantz. "Exploring User Perceptions of Discrimination in Online Targeted Advertising". In: *26th USENIX Security Symposium (USENIX Security 17)*. Vancouver, BC: USENIX Association, Aug. 2017, pp. 935–951. URL: https : / / www . usenix . org / conference / usenixsecurity17 / technical – sessions/presentation/plane.

[87] Isaac Prilleltensky. "Psychology and the status quo." In: *American Psychologist* 44.5 (1989), pp. 795–802. DOI: 10.1037/0003-066X.44.5.795.

[88] *Public Attitudes Towards Online Targeting*. Centre for Data Ethics and Innovation. 2020. URL: https: / / www . gov . uk / government / publications / cdei-review-of-online-targeting (visited on 02/2021).

[89] Irwin Reyes, Primal Wijesekera, Joel Reardon, Amit Elazari Bar On, Abbas Razaghpanah, Narseo Vallina-Rodriguez, and Serge Egelman. ""Won't Somebody Think of the Children?" Examining COPPA Compliance at Scale". In: *Proceedings on Privacy Enhancing Technologies* 2018.3 (2018), pp. 63–83. DOI: 10.1515/popets-2018-0021.

[90] Takahito Sakamoto and Masahiro Matsunaga. "After GDPR, Still Tracking or Not? Understanding Opt-Out States for Online Behavioral Advertising". In: *2019 IEEE Security and Privacy Workshops (SPW)*. 2019, pp. 92–99. DOI: 10.1109/SPW.2019.00027.

[91] William Samuelson and Richard Zeckhauser. "Status quo bias in decision making". In: *Journal of risk and uncertainty* 1.1 (Mar. 1988), pp. 7–59. DOI: 10. 1007/BF00055564.

[92] Deborah H Schenk. "Exploiting the Salience Bias in Designing Taxes". In: *Yale J. on Reg.* 28 (2011), p. 253. DOI: 10.2139/ssrn.1661322.

[93] *Share of global smartphone shipments by operating system from 2014 to 2023*. Statista. 2020. URL: https : / / www . statista . com / statistics / 272307 / market – share – forecast – for – smartphone – operating – systems/ (visited on 09/2020).

[94] Swapneel Sheth, Gail Kaiser, and Walid Maalej. "Us and Them: A Study of Privacy Requirements Across North America, Asia, and Europe". In: *Proceedings of the 36th International Conference on Software Engineering*. ICSE 2014. Hyderabad, India: ACM, 2014, pp. 859–870. DOI: 10.1145/2568225.2568244.

[95] Katie Shilton and Daniel Greene. "Linking Platforms, Practices, and Developer Ethics: Levers for Privacy Discourse in Mobile Application Development". In: *Journal of Business Ethics* 155.1 (Mar. 2019), pp. 131–146. DOI: 10.1007/s10551-017-3504-8.

[96] Laura Shipp and Jorge Blasco. "How private is your period?: A systematic analysis of menstrual app privacy policies". In: *Proceedings on Privacy Enhancing Technologies* 2020.4 (Oct. 2020), pp. 491–510. DOI: 10.2478/popets-2020-0083.

[97] Ashish Kumar Singh and Vidyasagar Potdar. "Blocking Online Advertising - A State of the Art". In: *Proceedings of the 2009 IEEE International Conference on Industrial Technology*. ICIT '09. USA: IEEE Computer Society, 2009, pp. 1–10. DOI: 10.1109/ICIT. 2009.4939739.

[98] *Software developer gender distribution worldwide*. Statista. 2020. URL: https://www.statista.com/ statistics / 1126823 / worldwide – developer – gender/ (visited on 06/2021).

[99] *Special Eurobarometer 431 "Data protection"*. Tech. rep. European Commission, 2015. URL: https://ec. europa . eu / commfrontoffice / publicopinion / archives / ebs / ebs _ 431 _ en . pdf (visited on 02/2021).

[100] Ryan Stevens, Clint Gibler, Jon Crussell, Jeremy Erickson, and Hao Chen. "Investigating User Privacy in Android Ad Libraries". In: *Workshop on Mobile Security Technologies (MoST)*. 2012. URL: http: //citeseerx.ist.psu.edu/viewdoc/summary? doi=10.1.1.298.7556.

[101] Mohammad Tahaei, Alisa Frik, and Kami Vaniea. "Privacy Champions in Software Teams: Understanding Their Motivations, Strategies, and Challenges". In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI '21. New York, NY, USA: ACM, 2021, pp. 1–15. DOI: 10.1145/ 3411764.3445768.

[102] Mohammad Tahaei and Kami Vaniea. ""Developers Are Responsible": What Ad Networks Tell Developers About Privacy". In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems Extended Abstracts*. CHI '21 Extended Abstracts. New York, NY, USA: ACM, 2021, pp. 1–12. DOI: 10.1145/3411763.3451805.

[103] Mohammad Tahaei and Kami Vaniea. "A Survey on Developer-Centred Security". In: *2019 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*. June 2019, pp. 129–138. DOI: 10.1109/EuroSPW.2019.00021.

[104] Mohammad Tahaei and Kami Vaniea. "Code-Level Dark Patterns: Exploring Ad Networks' Misleading Code Samples with Negative Consequences for Users". In: *What Can CHI Do About Dark Patterns? Workshop at CHI '21*. 2021, pp. 1–5. URL: http://hdl.handle.net/20.500.11820/ea71877b-4def-4c2c-aa45-e148122b4f36.

[105] Mohammad Tahaei, Kami Vaniea, Beznosov Konstantin, and Maria K. Wolters. "Security Notifications in Static Analysis Tools: Developers' Attitudes, Comprehension, and Ability to Act on Them". In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. CHI '21. New York, NY, USA: ACM, 2021, pp. 1–17. DOI: 10.1145/3411764.3445616.

[106] Mohammad Tahaei, Kami Vaniea, and Naomi Saphra. "Understanding Privacy-Related Questions on Stack Overflow". In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. CHI '20. Honolulu, HI, USA: ACM, 2020, pp. 1–14. DOI: 10.1145/3313831.3376768.

[107] *The State of Mobile in 2020*. App Annie. 2020. URL: https://www.appannie.com/en/insights/market-data/state-of-mobile-2020/ (visited on 09/2020).

[108] *The Value of Personalized Ads to a Thriving App Ecosystem*. Facebook. 2020. URL: https://developers.facebook.com/blog/post/2020/06/18/value-of-personalized-ads-thriving-app-ecosystem/ (visited on 02/2021).

[109] Janice Y. Tsai, Serge Egelman, Lorrie Cranor, and Alessandro Acquisti. "The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study". In: *Info. Sys. Research* 22.2 (June 2011), pp. 254–268. DOI: 10.1287/isre.1090.0260.

[110] Joseph Turow, Jennifer King, Chris Hoofnagle, Amy Bleakley, and Michael Hennessy. "Americans Reject Tailored Advertising and Three Activities That Enable It". In: (Sept. 2009). DOI: 10.2139/ssrn.1478214.

[111] Blase Ur, Pedro Giovanni Leon, Lorrie Faith Cranor, Richard Shay, and Yang Wang. "Smart, Useful, Scary, Creepy: Perceptions of Online Behavioral Advertising". In: *Proceedings of the Eighth Symposium on Usable Privacy and Security*. SOUPS '12. Washington, D.C.: ACM, 2012. DOI: 10.1145/2335356.2335362.

[112] Christine Utz, Martin Degeling, Sascha Fahl, Florian Schaub, and Thorsten Holz. "(Un)Informed Consent: Studying GDPR Consent Notices in the Field". In: *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*. CCS '19. London, United Kingdom: ACM, 2019, pp. 973–990. DOI: 10.1145/3319535.3354212.

[113] Max Van Kleek, Reuben Binns, Jun Zhao, Adam Slack, Sauyon Lee, Dean Ottewell, and Nigel Shadbolt. "X-Ray Refine: Supporting the Exploration and Refinement of Information Exposure Resulting from Smartphone Apps". In: *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. CHI '18. ACM, 2018, pp. 1–13. DOI: 10.1145/3173574.3173967.

[114] Max Van Kleek, Ilaria Liccardi, Reuben Binns, Jun Zhao, Daniel J. Weitzner, and Nigel Shadbolt. "Better the Devil You Know: Exposing the Data Sharing Practices of Smartphone Apps". In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. CHI '17. ACM, 2017, pp. 5208–5220. DOI: 10.1145/3025453.3025556.

[115] Haoyu Wang, Zhe Liu, Yao Guo, Xiangqun Chen, Miao Zhang, Guoai Xu, and Jason Hong. "An Explorative Study of the Mobile App Ecosystem from App Developers' Perspective". In: *Proceedings of the 26th International Conference on World Wide Web*. WWW '17. Perth, Australia: International World Wide Web Conferences Steering Committee, 2017, pp. 163–172. DOI: 10.1145/3038912.3052712.

[116] Ying Wang, Ebru Genc, and Gang Peng. "Aiming the Mobile Targets in a Cross-Cultural Context: Effects of Trust, Privacy Concerns, and Attitude". In: *International Journal of Human–Computer Interaction* 36.3 (2020), pp. 227–238. DOI: 10.1080/10447318.2019.1625571.

[117] Jay (Hyunjae) Yu and Brenda Cude. "'Hello, Mrs. Sarah Jones! We recommend this product!' Consumers' perceptions about personalized advertising: comparisons across advertisements delivered via three different types of media". In: *International Journal of Consumer Studies* 33.4 (2009), pp. 503–514. DOI: 10.1111/j.1470-6431.2009.00784.x.

[118] Sebastian Zimmeck, Peter Story, Daniel Smullen, Abhilasha Ravichander, Ziqi Wang, Joel Reidenberg, N. Cameron Russell, and Norman Sadeh. "MAPS: Scaling Privacy Compliance Analysis to a Million Apps". In: *Proceedings on Privacy Enhancing Technologies* 2019.3 (2019), pp. 66–86. DOI: 10.2478/popets-2019-0037.

Appendices

A Survey Instruments

A.1 Screening Survey

1. Please select the statement that best describes your primary role at your current or most recent job.

   • I'm not employed • Jobs NOT related to computer science, informatics, computer engineering, or related fields • Designing products (e.g. UI designer, interaction designer) • Developing software (e.g. programmer, developer, web developer, software engineer) • Testing software (e.g. tester, quality analyst, automation engineer) • Managing software development (e.g. project manager, IT manager, scrum master) • Privacy and/or security engineering (e.g. security engineer, privacy engineer, penetration tester, ethical hacker, cryptographer) • Others (please specify)

2. How many years of experience do you have in software development? (numbers only)

3. How many years have you worked in mobile app development, specifically? (numbers only)

4. How many mobile apps have you worked on in the last 3 years? (numbers only)

A.2 Main Survey

[After the participant read the participant information sheet and consent form, and agreed to participant in the study.]

1. How many mobile apps have you worked on in the last 3 years? (numbers only)

2. *[Scenario description.]* Imagine that you are a shareholder in a software development company. Together with a small team, you created a [personal finance management/gaming] app. The app will be published in Europe and the United States and is mainly targeted towards adults (above age of 18). To monetise the app, you have decided to use the "Acme" ad network to show ads to your users.

   The Acme ad network offers a step-by-step Assistant – a graphical user interface that provides various configuration choices for integrating ads into your [personal finance management/gaming] app. The Assistant asks the developer several questions and then provides ad network configuration code based on the answers that can be imported directly into an app with minimal additional coding required.

   The following are the 5 questions asked by Acme's Assistant, please answer them as if you were developing the [personal finance management/gaming] app.

   I Which ad formats are you integrating?

   • Banner: A basic ad format that appears at the top & bottom of the device screen. • Interstitial: full-page ads appear at natural breaks & transitions, such as level completion. Supports video content. • Rewarded Video: ads reward users for watching short videos and interacting with playable ads and surveys. Good for monetizing free-to-play users. Supports video content. • Native: customisable ad format that matches the look & feel of your app. Ads appear inline with app content. Supports video content.

   II What level of graphics do you want for your ads?

   • Ads with highest graphics quality. These ads will work best on newer phones with the latest operating systems. • Ads with moderate to low graphics quality. These ads will work on most phones.

   III Which platform are you integrating Acme ad network on?

   • Android • iOS • Unity • Windows Phone

   IV Select the type of ads that you want to show. [Participants were asked to choose between the personalized and non-personalized ads described according to the condition, to which they were randomly assigned. See the text of the options in section 3.1.]

   V Which of the following regulations apply to this app?

   • GDPR (General Data Protection Regulation) • CCPA (California Consumer Privacy Act) • COPPA (Children's Online Privacy Protection Act) • HIPAA (Health Insurance Portability and Accountability Act) • None of the above • I don't know

   VI What was the biggest reason that made you pick the ad type: [chosen ads type]? (Please provide as much as details you can. Your response helps us better understand the reasons behind your choices.) [Open-ended question]

   [Repeat the above questions for the second scenario.]

3. Assume that you decided to use personalized ads in both the gaming and financial management apps described earlier. How do you imagine you would go about asking for user consent for the personalized ads?

   • I'd use my own consent form • I'd use the consent form provided by the Acme ad network • I'd use a third-party consent form provided by a leading tech company (e.g., Facebook, Google, Amazon, Twitter) • I'd use a third-party consent form provided by a not-for-profit organization (e.g., Mozilla, Electronic Frontier Foundation) • I'd use a third-party consent form provided by other companies providing compliance services (e.g. OneTrust) • I won't ask for user consent because I don't think it's required • I won't ask for user consent because I don't think it's important • I won't ask for user consent because someone else in the team should take care of it • I won't ask for user consent because it's hard to do so • I won't ask for user consent because I'm not familiar with the consent process • I won't ask for user consent because the Acme ad network will take care of it • Other (please explain)

4. *[If "I'd use my own consent form" chosen.]* What information sources, if any, would you use to build your own consent form? [Open-ended question]

5. How, if at all, would your app's **revenue** change if you chose personalized ads over non-personalized ads in the **[personal financial management/gaming]** app described earlier? [Participants were asked about both app categories, in randomized order.]

- Decrease by more than 81% ● Decrease by 61%-80% ● Decrease by 41%-60% ● Decrease by 21%-40% ● Decrease by 1%-20% ● It won't change ● Increase by 1%-20% ● Increase by 21%-40% ● Increase by 41%-60% ● Increase by 61%-80% ● Increase by more than 81%

6. How, if at all, would the number of **users** of your app change if you chose personalized ads over non-personalized ads in the **[personal financial management/gaming]** app described earlier? [Participants were asked about both app categories, in randomized order.]

   ● Decrease by more than 81% ● Decrease by 61%-80% ● Decrease by 41%-60% ● Decrease by 21%-40% ● Decrease by 1%-20% ● It won't change ● Increase by 1%-20% ● Increase by 21%-40% ● Increase by 41%-60% ● Increase by 61%-80% ● Increase by more than 81%

7. How much priority do you give to privacy improvement and maintenance tasks in your daily development routines?

   ● Not a priority ● Low priority ● Medium priority ● High priority ● Essential

8. As a developer, how much control do **you** generally have over the amount of data collected by ad networks?

   ● No control at all ● Very little control ● Moderate control ● A lot of control ● Full control

9. How much control do **users** generally have over the amount of data collected by ad networks?

   ● No control at all ● Very little control ● Moderate control ● A lot of control ● Full control

10. What platforms have you previously developed apps for?

    ● Android ● iOS ● Blackberry ● Windows Phone

11. How involved have you been in in-app advertising activities? [Options were: Not at all, A little, A moderate amount, A lot, A great deal]

    ● Choosing an advertising partner or advertising network for an app. ● Configuring the types of in-app ads shown in an app (e.g., where to place ads, what categories of ads to show, etc.) ● Integrating the necessary code into an app to enable in-app advertising. ● Other (please specify)

12. Regarding mobile apps, have you used or worked with any advertising networks?

    ● AdColony ● Amazon Mobile Ad Network ● Facebook Audience Network ● Flurry ● Google AdMob ● InMobi ● Millennial media ● Twitter MoPub ● Unity Ads ● Vungle ● Greyfriars Bobby ● I have never included any ad networks in my mobile apps

13. *[If "I have never included any ad networks in my mobile apps" chosen.]* What are the primary reasons that you never included any ad networks in your apps? (Please provide at much as details you can. Your response helps us better understand your reasons behind your choices.)

14. What is the revenue model of the apps that you typically develop?

    ● Free with In-App Advertising, users cannot pay a fee to remove advertisements ● Free with In-App Advertising, users can pay a fee to remove advertisements ● Freemium model (app is free, certain features cost user's money) ● Paid download ● In-App purchases (selling physical or virtual goods through the app) ● Subscription (similar to Freemium, except instead of paying for extra features, users must pay for extra content) ● My apps are completely free ● Cannot remember ● Other (please specify)

15. Who decides what revenue model to use in the apps that you develop?

    ● Only me ● Developer(s) / Programmer(s) ● Project manager(s) ● CEO and/or other upper-level management ● Investor(s) ● Other (please specify) ● I do not know who was involved in the decision process

16. Who decides what advertisement network to use in the apps that you develop?

    ● Only me ● Developer(s) / Programmer(s) ● Project manager(s) ● CEO and/or other upper-level management ● Investor(s) ● Other (please specify) ● I do not know who was involved in the decision process

17. What is your main source of income in software or mobile development?

    ● I don't make money from software or mobile development ● Salary, not dependent on software/app revenue ● Primarily salary and bonuses, partially dependent on software/app revenue ● Primarily direct software/app revenue ● Other (please specify)

18. What type of employment best describes your most recent app development experience?

    ● Full time employee (or contractor equivalent) ● Part-time employee (or contractor equivalent) ● Freelance/consultant ● Furloughed (temporarily laid off) or on leave ● Unemployed ● Student ● Retired ● Other (please specify)

19. Please select the statement that best describes your primary roles at your most recent job.

    ● I'm not employed ● Jobs NOT related to computer science, informatics, computer engineering, or related fields ● Designing products (e.g. UI designer, interaction designer) ● Developing software (e.g. programmer, developer, web developer, software engineer) ● Testing software (e.g. tester, quality analyst, automation engineer) ● Managing software development (e.g. project manager, IT manager, scrum master) ● Privacy and/or security engineering (e.g. security engineer, privacy engineer, penetration tester, ethical hacker, cryptographer) ● Others

20. How many years of experience do you have in software development? (numbers only)

21. How many years have you worked in mobile app development specifically? (numbers only)

22. Where did you mainly learn to program and develop software?

    ● Self-taught ● High school courses ● College or university courses ● Online courses ● Industry or on-the-job training ● Others

23. How many people were employed in the organization for which you worked on the app development most recently?

    ● 1-9 employees ● 10-99 employees ● 100-999 employees ● 1,000-9,999 employees ● 10,000+ employees

24. How many members were in the team that you have directly worked with most recently? (numbers only)

25. How old are you? (numbers only)

26. In which country do you currently reside? [List of countries]

27. If you can't find your country in the above question options, please enter it here. [Open-ended question]

28. What is your gender?

   ● Male ● Female ● Non-binary ● Prefer not to say ● Prefer to self describe

29. If you'd like to be included in the raffle, please provide your email address.

30. Do you have comments or anything to say about the survey or study in general? (optional)

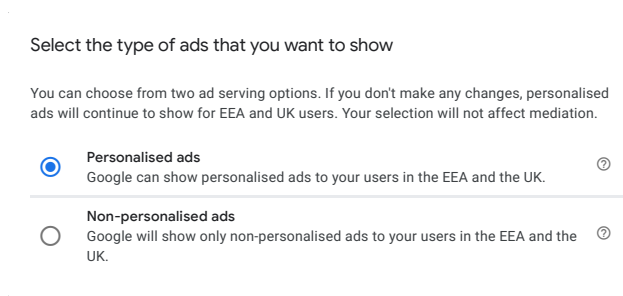## B Ads Personalization Options on Google AdMob Developer Dashboard

Figure 2: Screenshot from Google AdMob developer dashboard: Blocking controls > Manage EU user consent (Jan'21).
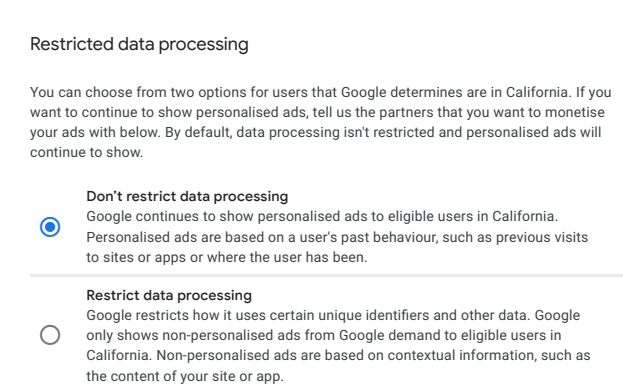

Figure 3: Screenshot of Google AdMob developer dashboard: Blocking controls > Manage CCPA settings (Jan'21).

C Participants' Demographics and Opinions About Ad Networks
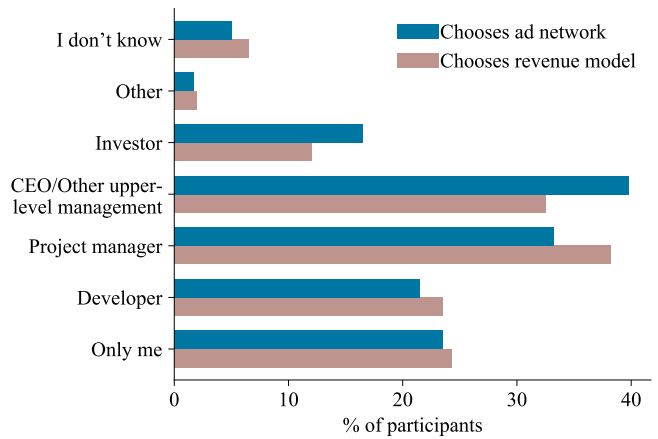

Figure 4: Responses about who decides what revenue model and ad network to use in the apps participants develop.
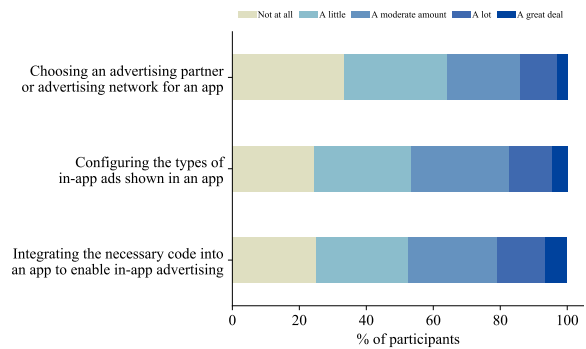

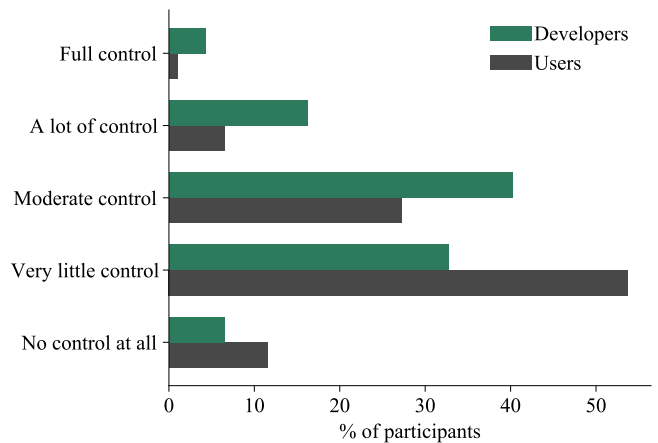Figure 5: Involvement in in-app advertising activities.

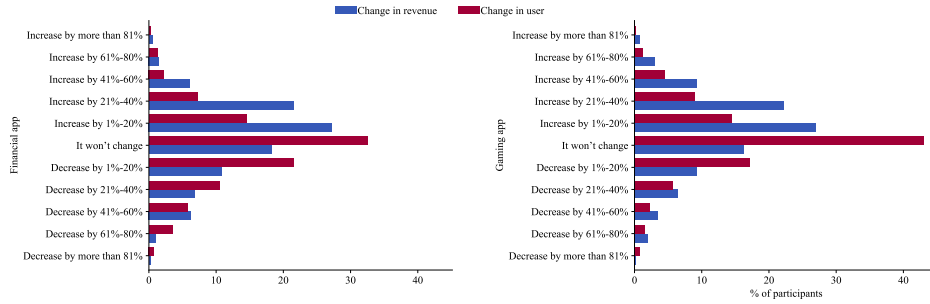Figure 6: Perceived control over ad networks' data collection.

Figure 7: Expected change in app's revenue and N of users if personalized ads are chosen over non-personalized ads.

	#Participants		#Participants
Age	$\mu = 27.4, \sigma = 8$	**Revenue Models**	
Gender		Free with In-App Advertising	120 (30.0%)
Male	330 (82%)	Completely free	103 (25.8%)
Female	58 (14%)	Freemium model	103 (25.8%)
Prefer not to say	11 (3%)	In-App purchases	83 (20.8%)
Non-binary	1 (<1%)	Free with In-App Advertising	82 (20.5%)
Current Continent of Residence		Subscription	54 (13.5%)
Europe	265 (66%)	Paid download	43 (10.8%)
North America	75 (19%)	Other	11 (2.8%)
Asia	24 (6%)	Can't remember	8 (2.0%)
Oceania	15 (4%)	**Which Ad Networks Used in the Past**	
South America	11 (3%)	Google AdMob	191 (47.8%)
Africa	7 (2%)	Never included any ad networks in apps	123 (30.8%)
Prefer not to say	3 (1%)	Facebook Audience Network	117 (29.2%)
Employment Status		Unity Ads	81 (20.2%)
Full-time	147 (37%)	Amazon Mobile Ad Network	64 (16.0%)
Student	107 (27%)	AdColony	33 (8.2%)
Freelance/consultant	75 (19%)	Twitter MoPub	27 (6.8%)
Part-time	54 (14%)	Flurry	15 (3.8%)
Unemployed	10 (2%)	InMobi	12 (3.0%)
Temporarily laid off	3 (1%)	Other	11 (2.8%)
Other	2 (<1%)	Vungle	9 (2.2%)
Retired	2 (<1%)	Millennial media	7 (1.8%)
Number of Employees		**Sources of User Consent Forms**	
1–9 employees	170 (42%)	The Acme ad network's form	128 (32.0%)
10–99 employees	142 (36%)	Leading tech company's form	90 (22.5%)
100–999 employees	49 (12%)	My own consent form (see Table 4)	71 (17.8%)
1,000–9,999 employees	21 (5%)	Not-for-profit organization's form	43 (10.8%)
10,000 or more employees	18 (4%)	Won't ask for user consent because:	39 (9.75%)
Team Members	$\mu = 7.3, \sigma = 10.3$	Acme ad network will take care of it	14 (3.5%)
Years of Experience		Someone else in the team should do it	14 (3.5%)
In software development	$\mu = 5.1, \sigma = 5.3$	Not familiar with the consent process	6 (1.5%)
In mobile development	$\mu = 2.7, \sigma = 2.6$	It's not important	2 (0.5%)
Number of Developed Apps in the Past Three Years	$\mu = 3.5, \sigma = 4.2$	It's hard to do so	2 (0.5%)
Software-Related Roles ($N = 291$)		It's not required	1 (0.2%)
Developing software	186 (64%)	Companies providing compliance services	24 (6.0%)
Testing software	37 (13%)	Other	5 (1.2%)
Managing software development	32 (11%)	**Given Priority to Privacy in Development Routines**	
Designing products	30 (10%)	High priority	144 (36%)
Privacy & security engineering	5 (2%)	Medium priority	136 (34%)
Main Income Source		Essential	61 (15%)
Salary, not dependent on software/app revenue	172 (43%)	Low priority	54 (14%)
Salary, partially dependent on software/app revenue	85 (21%)	Not a priority	5 (1%)
I don't make money from software/app dev.	80 (20%)	**Where Learned to Develop Software**	
Direct software/app revenue	58 (14%)	Self-taught	248 (62.0%)
Other	5 (1%)	College or university courses	237 (59.2%)
		Online courses	170 (42.5%)
		Industry or on-the-job training	103 (25.8%)
		High school courses	70 (17.5%)
		Other	3 (0.8%)

Table 5: Summary of participants' demographics and prior experience with ads ($N = 400$, unless otherwise specified).